

Contact Sensors for Dexterous Robotic Hands

by

David Mark Siegel
B.S.E. Electrical Engineering and Computer Science
Princeton University
(1983)

©Massachusetts Institute of Technology 1986

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the System Development Foundation, in part by the Office of Naval Research under contract N00014-81-0494, and in part by the Advanced Research Project Agency under Office of Naval Research contracts N00014-85-K-0124 and N00014-82-K-0334. Support for the author is provided in part by a National Science Foundation fellowship.

To my Mother, my Father, my Sister Sharon,
and my Grandmother Helen

**Contact Sensors
for
Dexterous Robotic Hands**

by

David Mark Siegel

*Submitted to the Department of Electrical Engineering and Computer Science
on April 23, 1986 in partial fulfillment of the requirements for the
Degree of Master of Science in Electrical Engineering and Computer Science*

Abstract. This thesis examines a tactile sensor and a thermal sensor for use with the Utah-MIT dexterous four fingered hand. Sensory feedback is critical for full utilization of its advanced manipulatory capabilities. The hand itself provides tendon tensions and joint angles information. However, planned control algorithms require more information than these sources can provide.

The tactile sensor utilizes capacitive transduction with a novel design based entirely on silicone elastomers. It provides an 8×8 array of force cells with 1.9 mm center-to-center spacing. A pressure resolution of 8 significant bits is available over a 0 to 200 grams per square mm range. The thermal sensor measures a material's heat conductivity by radiating heat into an object and measuring the resulting temperature variations. This sensor has a 4×4 array of temperature cells with 3.5 mm center-to-center spacing. Experiments show that the thermal sensor can discriminate among material by detecting differences in their thermal conduction properties. Both sensors meet the stringent mounting requirements posed by the Utah-MIT hand. Combining them together to form a sensor with both tactile and thermal capabilities will ultimately be possible.

The computational requirements for controlling a sensor equipped dexterous hand are severe. Conventional single processor computers do not provide adequate performance. To overcome these difficulties, a computational architecture based on interconnecting high performance microcomputers and a set of software primitives tailored for sensor driven control has been proposed. The system has been implemented and tested on the Utah-MIT hand.

The hand, equipped with tactile and thermal sensors and controlled by its computational architecture, is one of the most advanced robotic manipulatory devices available worldwide. Other ongoing projects will exploit these tools and allow the hand to perform tasks that exceed the capabilities of current generation robots.

Thesis Supervisor: Dr. John M. Hollerbach
Associate Professor of Brain and Cognitive Sciences

Acknowledgments

I'd like to thank my advisor John Hollerbach for help, encouragement, and support throughout this project. John's generosity is greatly appreciated; without the laboratory and manpower resources that he provided this project could not have been undertaken.

The tactile sensors described in this thesis would not have been possible without the tremendous effort made by Iñaki Garabieta, perhaps the finest craftsman one can find. Iñaki was a partner in this project; he was instrumental in the development of the sensing devices. There were times when silicone rubbers were out to get us; none worked quite the way we wanted. But in the end Iñaki could not be discouraged; he would not give up. And the silicone rubbers finally cooperated!

Sundar Narasimhan and I spend many hours hacking away on the computational architecture partially described in this thesis. I've spent more time working with Sundar over the past three years than with anyone else, and greatly appreciate all the help he has given during this project.

Luckily Harry Voorhees' tolerates my sense of humor, and can laugh at the same story several hundred times. The graphs in this thesis were made possible by his elegant and artificially intelligent plotting package. I even tried his copy-array

function once; it is truly impressive.

Laurel Simmons helped out with the thermal sensor and read a draft of this thesis. I'm grateful for her assistance. Even more importantly, Laurel helped make my experiences at the Lab more pleasant. For that I'm really thankful.

I'd also like to thank Scott Jones and Peggy Fong for keeping me well fed. Scott has also provided invaluable help by reading drafts of papers leading to this thesis, and by keeping the laboratory's printers running.

Chris Atkeson, Ki Suh, and Chae An have entertained me with the macho Asada arm, Jerry Roylance let me pester him with analog circuits, and Chris Lindblad poked around with most of the computers around the Lab, keeping them running.

Noble Larson, John Cox, and John Purbrick helped me out along the way. Thanks!

Ron Wiken provided great help keeping the parts room in shape, and he sure came up with a clever way for keeping the drawers from falling out of the resistor kit.

Klaus Biggers and George Gerphedie, both from the University of Utah, provided much help along the way. George spent a summer working on the computational architecture described in this theses. Dave Kriegman also helped out with the computer development. He was the first one to hack on the DMA code, and for that he deserves some kind of prize, maybe a leather bound copy of the code?

There are a number of people who have nothing to do with MIT, but who none the less provided support, encouragement, and kindness that kept me going. Josh Hinerfeld, formally of the Boston Housing Authority, has been a friend for years. Laura Metsch stumbled into Boston one summer, and carefully listened to lengthy discussions on laboratory space allocation issues. Bonnie Steinberg has greatly increased my understanding of biochemistry over the years. Pierre duPont taught me how to eat more pizza in one meal than I ever thought possible. Jay Hutter sure knows how to pick VCR movies. Hilary Lewis has entertained me to no end just by coming to MIT.

Contents

1	Contact Sensing and Dexterous Hands	1
1.1	Relation to Other Sensing Modalities	5
1.2	A Sensor Design Methodology	6
1.3	The Utah-MIT Hand	8
1.4	Scope of this Thesis	10
2	Tactile Sensing	13
2.1	Overview	13
2.2	Previous Work	17
2.2.1	Conductive Elastomer Sensors	17
2.2.2	Solid State Sensors	19
2.2.3	Piezoelectric Sensors	20
2.2.4	Magnetic Based Sensors	20
2.2.5	Optical Sensors	21
2.2.6	Capacitive Sensors	22
2.3	Design Issues	22
2.3.1	Transduction Principles	23

2.3.2	Fabrication Details	23
2.3.3	Detection Electronics	30
2.3.4	Design Process Summary	39
2.4	Performance Results	40
2.4.1	Quantifying a Sensor's Performance	40
2.4.2	Pressure Sensitivity	43
2.4.3	Hysteresis	45
2.4.4	Measurement Repeatability	47
2.4.5	Spatial Selectivity	47
2.4.6	Shape Discrimination	49
2.4.7	Interference Immunity	60
2.5	Overall Device Evaluation	60
3	Thermal Sensing	63
3.1	Robotics Applications for Thermal Sensing	64
3.2	Design of the Thermal Sensor	65
3.2.1	Theoretical Analysis of Sensor Performance	66
3.2.2	Performance Issues	71
3.2.3	Fabrication Details	72
3.2.4	Detection Electronics	75
3.3	Experimental Result	78
3.3.1	Material Identification	78
3.3.2	Measurement Repeatability	80
3.3.3	Variation in Heat Output	81
3.3.4	Thermal Temperature Recovery	82
3.3.5	Pressure Effects	83
3.3.6	Spatial Selectivity	85
3.4	Combining the Tactile Sensor and the Thermal Sensor	85
3.4.1	Performance Issues	85
3.4.2	Potential Sensor Construction	87

4	A Computational Architecture for Sensor Control	89
4.1	Design Methodology	90
4.2	Controller Architecture	92
4.3	Software Components	95
4.3.1	Servo Loop Scheduling	96
4.3.2	Intertask Communications	101
4.4	Performance Evaluation	107
5	Conclusions	111
5.1	Where Have We Gone So Far?	112
5.2	What Comes Next?	113
5.2.1	Improving the Tactile and Thermal Sensors	113
5.2.2	Sensor Based Hand Motions	115

List of Figures

1.1	The Utah-MIT Hand	8
1.2	The Utah-MIT Hand Positioning Table	9
2.1	Elastic-Dielectric Material	26
2.2	Dual Material Dielectric Layer	27
2.3	Tactile Sensor Cross Section	28
2.4	Prototype Tactile Sensor Packaging	29
2.5	Photographs of Tactile Sensor Components	30
2.6	Capacitance Array Schematic Diagram	31
2.7	Methods for Detecting Capacitance	33
2.8	Capacitance Array Scanning Electronics	36
2.9	Detailed Model of Tactile Scanning Electronics	37
2.10	Tactile Electronics Crosstalk Rejection	38
2.11	Tactile Sensor Tester	42
2.12	Tactile Pressure Sensitivity	44
2.13	Tactile Hysteresis	46
2.14	Tactile Spatial Selectivity	48

2.15 Tactile Image: Edges	51
2.16 Tactile Image: Edges	52
2.17 Tactile Image: Edges	53
2.18 Tactile Image: IC Sockets	54
2.19 Tactile Image: IC Sockets	55
2.20 Tactile Image: Two Edged Part	56
2.21 Tactile Image: Two Edged Part	57
2.22 Tactile Image: Bolt and Washer	58
2.23 Tactile Image: Molex Connector	59
3.1 Thermal Sensor Model	68
3.2 Behavior of Thermal Sensor Governing Equation	70
3.3 Thermal Sensor Printed Circuit Board Layout	73
3.4 Photographs of the Thermal Sensor	74
3.5 Thermal Sensor Cross Section	75
3.6 Thermal Array Scanning Electronics	76
3.7 Thermal Response to Various Materials	79
3.8 Thermal Steady State Response	80
3.9 Thermal Repeatability	81
3.10 Thermal Variations in Heat Output	82
3.11 Thermal Temperature Recovery	83
3.12 Thermal Pressure Response Variation	84
3.13 Thermal Spatial Selectivity	86
3.14 Tactile and Thermal Sensor Cross Section	88
4.1 Microprocessor Controller Block Diagram	93
4.2 Sample Use of the Servo Loop Scheduler	97
4.3 Servo Loop Scheduler Timing Diagram	100
4.4 Message Passing System Components	104
4.5 Typical Message Passing System Specification	106

List of Tables

2.1	Research and Commercial Tactile Sensors	18
2.2	Tactile Sensor Calibration Data	50
3.1	Thermal Properties of Materials	66
4.1	Computational Components of Hand Controller	92
4.2	Message Buffer Fields	102
4.3	Routing Table Example	103

Contact Sensing and Dexterous Hands

Human hands have amazing versatility. Think of the wide variety of tasks that we effortlessly perform with them each day. They are used for manipulating objects, static grasping, and exploratory motions. Our hands are so effective at what they do that we often take them for granted. Take the task of pulling change out of a pocket. Since visual information is of no use for this operation, our hand's tactile capabilities play a crucial role in locating a coin. Then our fingers readily grasp the coin and perhaps re-orient it for better stability. Thermal sensing capabilities can also play a role by distinguishing the metal of the change from other objects that might be in the pocket. Finally, the soft surfaces of our fingertips helps prevent the coin from slipping.

Our hands perform so many tasks with ease that they must have a particularly effective design. Even more amazingly, they operate in entirely unstructured environments. When we are mending clothing no special jigs for accurate placement of stitches are required. Instead, our hands pick up the fabric and start sewing. Our fingers can position the material and move around the needle without any assistance.

In some cases our hands are not capable of performing a task alone and tools

are required to help out. But think of how easily hands can use tools. It is hard to imagine a more versatile manipulatory system than a pair of hands with a suitable set of tools. In fact, the power of our hands and their ability to operate tools may very well be responsible for our success as a species.

Today's robotic end effectors, on the other hand, cannot rival the performance of our hands. One of their biggest deficiencies is that they must operate in structured environments. It is assumed that the objects they manipulate are well modeled and that their positions are accurately known. When the structuring assumptions that have been made fail, the robot is often helpless and cannot perform the desired task correctly. They usually lack the kinds of sensory feedback necessary for analyzing the world and modifying their actions to compensate for problems.

In addition to a lack of sensing capabilities, their mechanical dexterity is also not adequate for operating in an unstructured environment. Special jigs are required to make up for a robot's lack of flexibility. The cost involved with making a robot's workspace suitable for it to operate in often exceeds the cost of the robot itself. Many manufacturing and assembly operations cannot be automated because of these severe limitations. More importantly, some uses for robots are entirely precluded by their need for a structured environment.

Before robotic hands suitable for operation in unstructured environments can be made, advances in manipulators, actuators, sensors, computer controllers, control laws, planners, and robot programming languages are needed. Manipulators must be made more dexterous to permit complex grasping and orientation strategies. Actuators must be made faster and more compact. Sensors must be durable, and should provide information that insure the robot correctly performs its task. Computers need to be more powerful to handle the increased mechanical dexterity and the wealth of sensory data. Control laws must be extended to handle the mechanical structure of the advanced manipulators. Finally, tools must be developed to reduce the difficulty of programming a complex robot.

Investigation of grasping and manipulatory strategies for a sensor equipped hand are needed before human-like performance can be achieved. One of the many open questions in hand control is how to determine the most effective grasp orientation for a particular object geometry. The definition of "most effective" depends

on the particular operation intended; for some tasks a *power grip* that imparts maximum holding force onto an object is best. Other times, a *mobility grip* that gives an object its greatest range of motions would be preferred.

Contact sensors provide information that can simplify grasping operations. One important use for this information is to verify a grip's stability. If a grip is failing, contact sensors can provide information on the direction the object is slipping, possibly allowing corrective action to be taken. If the angle of contact between a finger and an object is approaching the limits of the contact frictional cone, indicating that an unstable situation may be imminent, the finger orientations can be modified.

Delicate operations also require sensory feedback. To avoid breaking an egg, for example, the finger's contact force must be small, yet large enough to avoid slippage. Programming a hand to write on paper probably requires carefully monitoring the position of the pen with respect to the fingers and paper. Flipping the pages of a book is tricky as well; a light grip on the paper is required to avoid ripping the page, while enough contact force must be applied to generate the friction needed to flip the sheet upward. These operations are even hard for humans, as we all know from experience.¹

The rewards for building an advanced dexterous robot that rivals human performance are great. Assembly operations could be more highly automated. Special purpose part feeders might be replaced by a more general purpose hand. Contact sensors could resolve a part's location and orientation, allowing the assembly operation to proceed smoothly. Thermal sensors could assist in selecting a particular object from a bin of parts. For insertion operations the tactile system could detect jams, and then direct the hand to take the appropriate corrective action. A robot polishing a material's surface could use texture detection to judge the quality of the finish.

A robot end effector that can operate in an unstructured environment will do more than advance the state of industrial automation, it can also relieve humans

¹Humans, of course, learn from their mistakes. Ideally, a hand should also benefit from its failures. Contact sensors can also provide important information for automatically diagnosing why the planned operation failed, and what future corrective action would be best to take.

of tasks that they should not be performing. Undersea exploration, nuclear power plant repair, space exploration, and rescue operations are hazardous and unpleasant for humans to perform. An advanced dexterous robot that can operate in an unstructured environment could be put to good use in these areas.

A hand can be used for more than just grasping and manipulating objects. Haptics, where the hand is used as an exploratory tool, is another important application. Exploratory motion is highly sensory driven; both contact sensors and joint positions provide information about an object and its environment. Texture and thermal clues can help a robot identify a material. Tracking an edge or seam to ascertain an object's outline can be done with tactile feedback. Detecting the orientation of a small part using its force contact profile is possible. Identifying a grasped object by mapping the finger contact locations and material type at those contact points to a library model has already been theoretically studied by Grimson and Lozano-Perez [1985]. The use of a hand as a sensory organ allows a robot to explore its environment and gather information that will help it perform its manipulatory functions successfully.

By building and programming a robotic device with performance capabilities similar to that of a hand, we can better understand biological motor control. Instead of just speculating on how the brain controls our hands, and on how sensory data is utilized, we can actually test theories on the robotic device. We can also go the other direction. That is, by studying human motion strategies we can develop better robotics control techniques. A manipulator with human-like mechanical and sensor capabilities allows us to use the every day experience we have using our own hands when programming the robotic hand.

This thesis concentrates on two of the aforementioned problems: the lack of adequate sensors and the need for high performance computer controllers. In the next section, an overview of some basic sensor driven control tasks that a dexterous hand could perform is presented. The introduction then concludes with an overview of the work covered in the remainder of this thesis.

1.1 Relation to Other Sensing Modalities

Sensors can be classified within the two broad categories of “non-contact” and “contact” devices. As one might ascertain from their names, non-contact sensors extract information passively, without making contact with the object being observed. Contact sensors, however, gather information by probing, poking, and otherwise disturbing the object at hand.

Vision is perhaps the most important non-contact sense that humans, and some robots, utilize. Sight gives a wealth of information about aspects of an environment that are unavailable to other sensory modalities. Uses for robotic vision systems include parts inspection, object identification, and character recognition.

It should be noted that machine vision can provide some, but not all, of the functionality of a contact sensing device. Thermal properties other than an object’s temperature can only be obtained by a contact sensor. While vision can give useful views of an object provided adequate lighting is available, a contact sensor can provide some of the same information in the dark. An object’s surface profile and texture can only be indirectly obtained from vision, with variations in lighting conditions and viewing angles affecting the results. Clearly, a contact sensor obtains this information more accurately.

Various other non-contact sensations are useful at times. Sound detection provides information in assembly operations. The “click” of two fastener coming together is a good indication that the task was completed successfully. Sonar is often used to help mobile robots navigate through rooms, and to help bats navigate through caves.

Another non-contact sense that most humans have, and some robots may have in the near future, is smell. Smells add not only to our taste of foods but they provide warnings of hazards that a visual system could not detect itself. A robotic *nose*, developed by Ikegami and Kaneyasu [1985], duplicates some of the human’s abilities, and can be put to use in situations where a human would rather not be sniffing.

Fusion of data from multiple sensor sources supplemented by a knowledge base describing world objects will ultimately be possible. Using information contributed

by different sensing modalities reduces the chance of error, increasing the robustness of the overall system. Humans have a very robust integrated and intelligent sensor system. We receive information from touch, sound, sight, and smell. The brain uses clues extracted from these sources, along with general knowledge of our world, to make all its sensor related decisions.

Sensors play an important role in the development of a dexterous robot capable of operating in an unstructured environment. The next section discusses how to design them.

1.2 A Sensor Design Methodology

The design of contact sensors should not be done in an ad hoc fashion. In particular, a number of questions should be answered before undertaking actual development. The following paragraphs pose these questions and then discuss their relevance to the sensor design process.

What do you want to sense? This question is not as obvious as it might seem. In many cases the answer depends on the sensor's intended uses. Tactile sensors, for example, mean different things to different people. For one particular application an array of binary contact detectors might suffice. Other applications might require 16 bits of force data, or shear detection in addition to surface normal force information. Slip detection is often cited as an important capability for sensors. Without a clear idea of the sensor's intended function, it is very hard to build a useful device.

How do you go about sensing this quantity? Before a sensor can be designed, there must be a clear understanding of the phenomenon being sensed, and how it manifests itself on the environment. In the case of slip detection, for example, it is not obvious what a slipping object does to a gripper that would indicate that it is really slipping. It is important to notice that this question should be answered separately from the question concerning the transduction technology to be employed. For example, to detect contact profiles a sensor must measure how its surface deflects when pressed against an object. So, contact profiles are sensed by measuring a surface deformation. This observation does not have any bearing

on what particular transduction technology should be used.

Where is the sensor to be used? The selection of an actual transduction process must take into account such practical issues as how much space is available for mounting the device. Equally important is to ascertain the expected operating environment. For example, will the device need special protection against damage? Special issues covering the packaging should also be addressed. A tactile sensor, for example, might need a compliant covering to increase frictional contact with a grasped object and for protection. Notice that these requirements can interact with the transduction mechanism. Providing a compliant covering might reduce force sensing resolution, for example.

What transduction processes can be used? The answers to previous questions guide in answering this question. For example, if the sensor measures force by inducing a positional change in a material, the transduction process must detect this positional change. If there is no transduction process capable of detecting the environmental disturbance of interest, the previous question concerning what one wants to sense will need to be re-examined.

How does the selected transduction process work? A full understanding of the transduction process will allow the sensor's performance to be optimized. In addition, flaws with the approach, if any, can be detected. The model of the transduction process should incorporate the sensor's interaction with its environment.

How should the sensor be fabricated? The answer to this question must incorporate the answers to all the previous questions. The ultimate fabrication techniques chosen should provide packaging suitable for both the transduction principles employed and for the intended use of the device; a tactile sensor that passed all performance specification with flying colors is of little use if it is too big to fit in its intended mounting location.

To begin the design process for the tactile and thermal sensors described within this report, a discussion of the Utah-MIT hand is necessary. This provides an answer to one of the most important of the previous questions: where the sensors are to be mounted.




Figure 1.1: *Photograph of The Utah-MIT Hand.*

1.3 The Utah-MIT Hand

The Center for Engineering Design at the University of Utah and the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology are developing a multi-fingered robot hand to pursue the achievement of advanced robot dexterity [Jacobsen et al. 1984]. The Utah-MIT hand has been made of approximate anthropomorphic size and design, partly because the human hand is obviously a particularly effective design, and partly because existing experience in using our own hands can aid in thinking of control strategies for the robot hand. The hand has four fingers, arranged as three fingers in a planar sequence and an opposing thumb. Each finger has four degrees of freedom, similar to a human finger save that the two degree of freedom knuckle joint is split into separate abduction/adduction and flexion/extension joints for technical reasons of tendon routing. Four fingers were chosen as a compromise to provide more functional capability than a three fingered hand [Mason and Salisbury 1985; Okada 1979] because of object reorientation ability and hand shaping functions, yet to avoid the added mechanical




Figure 1.2: The Utah-MIT hand positioning table.

complexity of five fingers.

Each finger is driven by 8 tendons, which is 3 more than required by the $n + 1$ minimum rule for tendons versus degrees of freedom [Mason and Salisbury 1985]. The transmission and actuation systems are two of the more novel aspects of this hand design. The tendons are composed of specially woven kevlar tapes, routed by pulleys across the finger joints and the wrist to the forearm actuator assembly. The 32 electropneumatic actuators have favorable power to weight ratios and can be adequately controlled for force. Sensors are provided for tendon tension and joint positions.

The hand itself is mounted onto the z axis of a 3 degree-of-freedom positioning table (Figure 1.2). The table's surface moves in the xy plane, while the z axis is translated by an independent y axis. The redundant y motion permits experimentation with tracking operations; an object moving in the y direction of the positioning surface can be independently tracked by the hand.

1.4 Scope of this Thesis

This thesis is motivated by the need to equip the Utah-MIT dexterous hand with contact sensors. Initially, an examination of commercially produced sensors produced disappointing findings. No devices suitable for mounting on the fingers of the hand were found. Research sensors were more promising, but there were still no refined technologies that could be directly used in this project. In fact, the performance requirements we set for the tactile sensor were not all that stringent; it was the space and mounting constraints, as well as reliability issues, that became limiting factors. Thermal sensors were in a much less developed state. Only at the completion of this project did we discover another thermal sensing effort, conducted concurrently and independently of this work [Russell 1985], that examined some of the same issues covered in this thesis.

The tactile sensor is based on variations in the distance between two parallel plates of a capacitor [Siegel et al. 1985]. As force on a sensor point is increased, the gap between the plates of a capacitor will decrease. The measured capacitance is related to the distance between the plates, and hence the force being applied at that point. A 8×8 array of capacitor cells is used, giving the device 64 force sensitive tactile points.

Thermal sensors can provide useful information about a manipulator's environment. These devices are designed to measure the properties of a material that are related to its heat conduction and heat capacity, as well as its absolute temperature. This sensing modality has received little attention in the literature, yet it can provide information that cannot be obtained from vision or tactile sensors alone. Humans, for example, have little trouble distinguishing metal from plastic. The great difference in their thermal conduction characteristics give the metal a distinctively cooler feeling.

The thermal sensor measures heat conduction by actively applying heat to an object, and measuring the temperature changes at the sensor's surface [Siegel and Simmons 1985]. If the material the sensor is in contact with is a good heat conductor, the surface temperature of the sensor will rapidly decline. This device uses surface mounted thermistors for temperature measurement and has a 4×4

array of thermal detection cells.

Since the potential usefulness of sensors is heavily dependent on their reliability, sensitivity, and packaging, much attention is given to these considerations. A detailed set of performance experiments was conducted for both the tactile and thermal prototypes to fully quantify their capabilities. The motivation for the particular experiments chosen is carefully discussed. In current literature, sensor performance is commonly characterized in an ad hoc fashion, making comparisons difficult.

The performance experiments conducted show that the sensors can be used in a number of useful ways. The tactile device's spatial resolution is sufficient for identification of small objects by recording their contact profiles. Objects larger than the sensor can be characterized by a groping strategy. The tactile sensor can classify the contact type between an object and its surface; discrimination between point, line, and plane contact is possible. The thermal sensor can identify materials in contact with its surface.

A sensor equipped hand alone cannot perform useful operations. High performance computers are required for the computationally complex sensor-based control strategies envisioned. This thesis also discusses a multiprocessor based computational architecture suitable for these needs. The system is based on high performance single board computers interconnected via a shared memory bus. Commercial products were chosen over custom designed hardware to reduce the system's development time.

To summarize, the important issues to be investigated in this project include:

- **Packaging:** most devices do not have suitable packaging for installation on the Utah-MIT hand's fingers. Many sensors are too large to fit within the tight space constraints at the mounting sites, and cannot bend around the curved surfaces of the fingers.
- **Reliability:** many sensors lack the reliability needed for repeated use in actual applications. Much attention is given to construction techniques that insure a long mean time between failures.
- **Thermal capabilities:** almost all prototype sensors, and all commercial prod-

ucts, cannot sense the thermal properties of a material. Further investigation was warranted to determine this sensing modality's potential usefulness.

- **Benchmark testing:** many research and commercial devices have not been adequately tested making it hard to determine their actual performance. By using a sensor test facility, detailed measurements are obtained, and the sensor's long term reliability established.
- **Computational support:** high performance computer hardware and appropriate software support must be provided for the sensor equipped hand. An architecture that meets the long term needs of this project has been designed and tested.

These issues are fully discussed in subsequent chapters of this thesis. Finally, concluding remarks on work in progress to fully utilize the tools described is presented. The Utah-MIT hand equipped with contact sensors and a powerful computational architecture will be used to perform complex manipulatory operations that robots have never before been able to achieve.

Tactile Sensing

The work presented within this chapter describes a tactile sensor for use with the Utah-MIT four fingered dexterous hand. Much of the development effort centered on isolating durable materials, and finding the construction and packaging techniques that permit mounting the sensor on the hand and reliably using it over an extended period of time. The following sections review past work in tactile sensing and describe the prototype's design and performance in depth.

2.1 Overview

The potential usefulness of tactile sensors in robotics control applications is well documented in the literature [Overton and Williams 1981; Harmon 1983; Dario and De Rossi 1985]. However, the number of robotic systems equipped with tactile sensors is surprisingly low. In fact, investigations of how to process and effectively utilize such information are equally uncommon. The slow progress in this area is generally attributed to the lack of sensors that have the spatial resolution, sensitivity, repeatability, durability, and packaging required for robotics applications. The last two of these criterion create the most important and difficult problems in sensor design, and are often ignored.

In addition, the lack of suitable dexterous robots has slowed progress in tactile sensing applications. To fully exploit the power of a tactile sensor a manipulator must have the ability to reposition its fingers in three dimensions, as does the Utah-MIT hand. The tactile sensor, of course, must be suitable for use with the intended robot; if the sensor is too large or cannot conform to the mounting surface available, it is of no use.

Commonly cited applications for tactile sensing include parts orientation, adaptive grasping, shape recognition, slip detection, and texture discrimination [Harmon 1982]. Each of these tasks is difficult, if not impossible, to accomplish without the aid of touch feedback. Imagine performing a manipulatory task with your own hands covered with heavy gloves; the lack of tactile feedback in this situation severely hinders coordination. For these reasons, equipping the Utah-MIT hand with tactile sensors is considered essential.

Both Binford [1972] and Harmon [1982] have speculated on the performance needed for a robotics tactile system. Harmon's report presents the responses to a questionnaire on touch sensing. The respondents were asked what properties an ideal tactile sensor should have. Spatial resolution was suggested to be in the range of 1 – 2 mm. The transduction matrix should have between 5×10 and 10×20 elements. Sensitivity ranges of 5 – 10 grams per sensor cell were given. Hysteresis was said to be "intolerable," yet skin like rubber coverings were desirable.¹

Harmon estimates that a touch sensor for automotive assembly should have a 12×12 array of force sensitive elements, or cells, in a 2.5×2.5 cm square area, giving an overall spatial resolution in the range of 1.6 mm. Hysteresis should be less than 5 percent, and a scan rate of 100 Hertz should be achieved. Finally, he recommends a minimum force sensitivity of 40 grams per sensor cell.

These figures are often cited, though little attention has been given to their validity. In actuality, no suitable manipulator equipped with a touch sensor has been available to test them. However, some general comments on sensor requirements are warranted. To successfully evaluate the needs of a tactile sensor system, a detailed description of the planned tasks must be compiled. For example, if the

¹This is actually a contradiction. A later section will discuss the issue of hysteresis in more depth.

sensor is to be used to detect the outline of small parts, the spatial resolution must be suitably dense. On the other hand, if a sensor is to be used to detect slippage, spatial resolution is of less importance. Picking a scan rate of say, 100 Hertz, independent of application does not seem reasonable either. Very high scan rates probably are needed for slip detection, while low scan rates would be suitable for small parts recognition.

Since the body of literature on how to use tactile sensor data is so small, selection of the appropriate sensor requirements is difficult. As a compromise, it seems reasonable to make the device approximate a human's touch sensing capabilities, since our hand's sensor driven dexterous performance is so good. A brief review of human touch sensing capabilities is presented for this purpose.

Humans have four different tactile sensors, each with their own specialized functions [Hollerbach 1984]. The Merkel and Ruffini receptors have some static touch response, while the Meissner and Pacinian respond better to a changing stimulation. Two point discrimination is on the order of 1 mm, a reasonable goal to achieve for robotic devices. Texture detection is mostly an active process, where fingers are moved along a surface. In this case, the Merkel, Meissner and Pacinian all contribute information.

Each tactile receptor's response is controlled by both their structure and their placement within the skin [Phillips and Johnson 1981]. The Meissner's corpuscle responds to tangential stress, the Merkel responds to compressive stress independent of orientation, and the Ruffini respond to directional skin stretch. Certain stimuli are detected by their mechanical interaction between the skin and the corpuscle's response pattern. Merkel endings, for example, best detect the high stress patterns that an edge induces within the skin. From this we can conclude that the ultimate tactile sensor might require more than one transduction technology, along with an appropriate skin-like mounting medium.

After selecting the sensor's performance requirements, many engineering challenges must then be addressed. The space constraints involved with mounting a tactile sensor are often severe, and the environment under which it must operate is commonly harsh. The active surface must withstand repeated contact with materials, and probably requires a convenient method for replacement when it becomes

worn or damaged.² The electronic and other delicate portions of the sensor must be protected from the environment, while simultaneously performing their sensing function. Clearly, special attention must be given to the sensor's overall packaging for these requirements to be met.

Specifically, a number of design goals for the tactile sensor were established. These requirements come from both the special needs of the Utah-MIT hand, and from what are generally felt to be desirable capabilities for a tactile sensor. The most important specifications include:

- An 8×8 array with under 2 mm cell center to center spacing: This requirement is motivated in part by the mechanoreceptor spacing of the human skin and in part by engineering constraints. Human tactile resolution is found to be on the order 1 mm [Phillips and Johnson 1981], approximately the same as the actual mechanoreceptor spacing. To approximate human performance a sensor spacing close to this level is desired. However, engineering considerations dictate a slightly larger spacing, to permit fabrication of a working and reliable sensor in a reasonable time frame. The 2 mm range was selected as a compromise between these two factors.
- A sensing technology that permits mounting the device on the non-planar surfaces of the hand: Since the tactile sensor is intended for use with the Utah-MIT hand, and since the hand's fingers have curved surfaces, non-planar mounting is a critical design goal.
- A low profile package since space at the finger mounting sites is at a premium: The tactile sensor must not interfere with the overall dexterity of the hand. Bulky sensors would reduce the finger's effective workspace.
- Small space requirements for the wiring and electronics needed at the fingertip site. Again, the tactile sensor must fit within the small space available at the hand's finger sites.
- A compliant surface covering for increased prehension stability [Fearing and Hollerbach 1984]. Recent work in grasping has shown that the mechanical prop-

²The human skin is useful in these respects; it is both highly durable and regenerates when worn or damaged. Although at times it too may not hold up in the harsh environments that some manipulatory tasks require.

erties of the finger's skin play an important role in grasping operations. A soft material will give a plane contact, instead of the point contact that would result from a hard material. This yields greater grasp stability in part from the increased friction at the contact area. Brockett [1985] proposes using a rheological surface that molds around the object being grasped, again to increase a grasp's natural stability.

One potential requirement that has been omitted from this list is the ability for the sensor to detect tangential force components. The sensor designed can only respond to forces normal to its surface. This severe shortcoming is dictated by engineering considerations; adding tangential force detection was determined to be too complex to undertake at the current time. Unfortunately certain operations that tactile sensors should perform are harder to do using just normal force detection. For example, detecting slippage in the most general case requires comparison of the tangential force component with the normal force component made by contact point friction.

2.2 Previous Work

Many tactile sensor designs have been reported in the literature and a few have recently been made into commercial products [Dario and De Rossi 1985; Ogorek 1985; Harmon 1982]. However, none of the reported sensors meet the previously described requirements for the Utah-MIT hand. Problems with packaging, reliability, size, and sensitivity preclude their use.

All tactile sensors must convert an applied force into a measurable electric signal. The design choice made is in how that conversion process is to be done. The most common technologies employed are based on optics, resistance, magnetic, or capacitance. This section discusses several research and commercial sensors (see Table 2.1) and briefly describes their functionality.

2.2.1 Conductive Elastomer Sensors

Several experimental tactile sensors based on a material's resistive properties have been developed including those of Purbrick [1981], Overton and Williams [1981],

Table 2.1: *Summary of research and commercial tactile sensors.*

Year	Device	Size	Spacing	Technology
1981	Purbrick	8×8	2.5 mm	resistive
1982	Hillis	16×16	0.7 mm	resistive
1982	Raibert	6×3	1.6 mm	VLSI
1983	Barry Wright	16×16	2.5 mm	resistive
1983	Hackwood	7×7	2 mm	magnetic
1983	Nomura	15×15	2 mm	optical
1984	Boie	8×8	2.5 mm	capacitive
1984	Chun	8×8	2 mm	VLSI
1984	Dario	8×16	3 mm	piezoelectric
1984	Jacobsen	1 cell	20 mm	optical
1984	Lord Corp	10×16	1.8 mm	optical
1984	Schneider	24×25	0.4 mm	optical
1985	Ten Grotenhuis	16×8	1.5 mm	carbon fiber
1985	Petersen	3×3	1.6 mm	VLSI
1986	Grahn	3×4	1 mm	piezoelectric

and Hillis [1982]. The Purbrick sensor uses conductive strips placed at right angles to printed circuit board traces, where tactile cells are formed at each intersection. As pressure on the conductive strip increases, the contact resistance formed with the circuit board trace decreases. The device achieves a 3 mm center-to-center cell spacing. The Overton and Williams sensor, based on conductive rubber plugs, achieves a 2.5 mm cell spacing.

Hillis' device is based on anisotropically conductive silicone rubber. This material is electrically conductive along one planar axis. As pressure is applied to the top of the rubber, its contact with a bottom circuit board is increased. By measuring the resulting resistance, the applied pressure at a particular point in the array can be inferred. Hillis constructed a 16×16 array of sensor points, with 1 mm center to center spacing.

Barry Wright Corporation's commercially available Sensoflex tactile system has

an 16×16 array with 2.5 mm center to center cell spacing. The conductive elastomer technology employed within this sensor appears to be based on Purbrick's design.

A few sensors utilizing carbon fibers have been pursued [Larcombe 1976; Ten Grotenhuis and Moore 1985]. Two sets of parallel carbon fiber layers are placed perpendicular to each other. As pressure is applied on the grid, the junction resistance between the upper and lower fibers decrease.

Unfortunately, most of these sensors suffer from high hysteresis and low sensitivity. In addition, many have inadequate cell spacing, though the Hillis sensor has achieved one of the highest spacings to date.

2.2.2 Solid State Sensors

Solid state VLSI based designs allow small transduction cells to be integrated with the amplification and scanning electronics, and with low level data analysis hardware. Raibert and Tanner [1982] have developed a VLSI tactile sensing computer. The chip contains an array of force sensing cells, each with their own processor. Two-dimensional convolutions and serialization of the sensor data are performed directly on the device. However, the force sensing cells are not based fully on VLSI. Instead they rely on a compressive resistive material placed over the package to transduce force changes to a measurable signal.

The two most common fully VLSI based force transduction techniques utilize capacitive variations [Chun and Wise 1985] and piezoresistive changes [Wong and Van der Spiegel 1985]. These sensors have silicone structures that respond to external force stimuli. The latter principle measures the deflection of a small silicone structure with an embedded pressure sensitive resistor, or piezoresistor. With the former technique, a small diaphragm is constructed with an upper and lower plate. As force is applied, the diaphragm will contract, resulting in a change in capacitance between the plates proportional to the applied force.

Transensory Devices VLSI tactile sensor [Petersen et al. 1985], now commercially available, uses piezoresistive strain gauge based elements spaced approximately 1.6 mm apart, in a 3×3 array. Each cell has a normal force range of 0 to 900 grams. A small pedestal placed on each cell concentrates an applied force onto

the sensing site, improving the overall performance.

While VLSI based designs show great promise, their current use in tactile sensors is limited by packaging difficulties. The integrated devices themselves are fairly durable, but fragile bonding wires must be brought out from the chip. In addition, VLSI sensors are currently fabricated only on flat surfaces, precluding their use on the curved portions of the Utah-MIT hand's fingers.

2.2.3 Piezoelectric Sensors

Dario et al. [1984] describes a sensor based on polyvinylidene fluoride (PVF₂), a material that has both piezoelectric and pyroelectric properties. Pressing a thin film sheet of PVF₂ generates an electrical charge proportional to the applied force. In addition, the Dario sensor has a resistive based sensing layer to obtain static force measurements; the PVF₂ only responds to changes in the force stimulation. The pyroelectric properties of the material can be used to measure a material's thermal properties. The device obtains a 3 mm center-to-center cell spacing in an 8×16 array.

Grahn and Astle [1986] describe a novel sensor based on ultrasonic pulse-echo ranging. A sheet of piezoelectric film is placed under a layer of silicone rubber. Electrodes are placed on the film at discrete points. To detect the applied pressure at a location of the array, the piezoelectric film is excited. The vibration pulse is reflected off the silicone rubber surface, and its return is detected by the film. The elapsed time between the pulse and the echo is proportional to the thickness of the silicone rubber at that point. While this sensor shows promise, the device has yet to be made into a small array suitable for finger mounting.

2.2.4 Magnetic Based Sensors

A magnetic based sensor has been proposed that extracts shear force as well as surface normal force [Hackwood et al. 1983]. An array of magnetic dipoles is embedded in an elastic medium. This rubber covering is placed over an array of magnetic field sensors fabricated in a silicone substrate. A pure normal force would increase all the magnetic field sensors around the elastic medium's dipole evenly.

If a shear force were applied, the dipole would rotate and translate in the elastic medium. The rotation and translation can then be measured by the variation in magnetic field in each of the cell's sensors.

Sensors of this type are interesting because of their ability to detect shear force, as well as surface normal force information. Most of the sensors described in the literature are unable to do this. However, the Hackwood magnetic sensor has not worked out well; the magnetic effects being measured are quite small, and hard to detect. In addition, the sensor is overly sensitive to extraneous magnetic fields and metallic objects.

2.2.5 Optical Sensors

Jacobsen et al. [1984] discusses a fiber optical based tactile sensor that utilizes a birefringent material. As force deflects this material, its ability to pass polarized light is reduced. An optical sensor measures the returning light, and converts it to an electrical output proportional to force. Unfortunately, only single tactile cells have been constructed and it is not clear if the technology can be extended to form an array.

Two commercially produced optical sensor are now available. An early Lord Corporation device (discussed in Ogorek 1985, and in Dario 1985) with an optical transduction mechanism has an 8×8 array with 7.6 mm cell spacing. As pressure on the sensor is increased, the deflection of its surface reduces the intensity of a light source. Lord's latest optoelectronic sensor achieves 1.8 mm cell spacing with a 10×16 array. The entire pad can be scanned in 3 milliseconds and its force capacity is 3 pounds. Begej Laboratories [Begej 1984] makes a fiber optic sensor with 1.0 mm element spacing. The system images the tactile output by projecting the optical fiber output onto a solid state camera. Both these sensors are too bulky for mounting on the Utah-MIT hand.

Schneider and Sheridan [1984] describe a silicone rubber based optical sensor. A reflective coating is applied to the underside of a silicone rubber sheet. A bundle of fibers project light onto the sheet, and receive the reflected light back. Pressure applied to the silicone rubber modulates the reflected light's intensity. A video

camera is used to read the sensor's output by imaging the fiber bundle. The device achieves a very high spatial resolution of 330 points per square centimeter in an array of 34×35 sensor elements.

An advantage of optical sensor designs is their relative insensitivity to electrical interference. Since the force signal is propagated from the sensor site to detecting electronics optically, noise immunity is quite good.

2.2.6 Capacitive Sensors

A capacitive based sensor measures an applied force by detecting a change in the separation of two electrically conductive plates. As the plate separation is reduced, the capacitance between them increases.

Boie [1984] describes a prototype sensor with an 8×8 cell array and 2.5 mm center-to-center spacing. The technology described in his work is promising and was selected as a starting point for the design used within this report. The remaining sections of this chapter discuss capacitive based tactile sensing, and then describe the sensor design, construction, and performance in depth.

2.3 Design Issues

This section discusses the tactile sensor's design. The transduction principles, the fabrication processes, and the scanning electronics are covered in depth. Attention has been given to describe both the final sensor prototype and some intermediate devices. This approach gives a better feel for the difficulties encountered during the project, and better motivates the choices made in the final design.

The sensor described within this report is based on the work of Boie [1984]. Both sensors share the capacitive based force transduction principle, described in Boie's paper. This work, however, investigates the use of improved construction techniques to attain greater durability and to permit mounting on curved surfaces, uses different scanning electronics for improved performance and reduced size, has greater spatial resolution and sensitivity, and can be scanned at a higher rate.

2.3.1 Transduction Principles

Two parallel electrically conductive plates generate a capacitance proportional to their separation distance. If a compressible dielectric material is placed between them, a force applied to the capacitor's top will reduce the plate separation distance. The resulting change in capacitance can be used to infer the applied force. This principle is the basis of the capacitive tactile sensor's force transduction mechanism.

An array of capacitive cells is formed by sandwiching a dielectric layer between two sets of parallel conducting traces, with the top etches perpendicular to the bottom ones. A capacitor is formed each time an upper trace intersects a lower trace. To make an array of 64 force sensing capacitors, 8 upper traces and 8 lower traces are used.

2.3.2 Fabrication Details

As previously mentioned, the basic operating principle for the tactile sensors is the measurement of an applied pressure by detecting a variation in the gap of two parallel capacitive plates. Hence, the material between the two plates is a crucial component of the device; it forms both the elastic layer that compresses in response to pressure, and the dielectric layer that provides the capacitance between the two plates.

Ideally, this material should compress linearly as force on it increases. This spring-like behavior is given by

$$F = k\Delta X, \quad (2.1)$$

where F is the applied force, k is the effective spring constant, and ΔX is the positional change that the force produces. Having a material that closely approximates this equation makes it easy to translate the output of the sensor back to the applied force, since the detected capacitance will then vary in linear correspondence to the force. Thus, the material selected should have linear compression over the range of pressures to be applied.

The dielectric layer is composed of an electrically insulating silicone rubber. Silicone rubber has the desirable flexibility and durability that the layer requires.

In addition, its dielectric constant is approximately 4, which increases the overall performance of the device by increasing each cell's effective capacitance.

Although the silicon rubber is flexible, it does not have desirable continuum mechanical properties. In fact, it is largely incompressible when subjected to a uniform pressure applied over its surface. Since the tactile sensor detects an applied stimulus by measuring the resulting deformation profile that it generates, a solid object pushing uniformly into the pad will not excite a response. The following paragraphs examine this problem in more depth, and explain the solution that has been employed.

To better understand the behavior of the rubber medium, consider it to be composed of small unit volume cubes with xyz coordinate frames. Assume that the yz plane is aligned with the material's surface, and forces are applied along the x axis. The pressure applied to an axis of the cube is called a stress. A positional change resulting from the stress is called a strain. Stresses are related to strains by Hooke's law, where E is the modulus of elasticity.

For the small cube, Timoshenko and Goodier [1951] gives the relationship between stresses and strains as

$$\begin{aligned}\epsilon_x &= \frac{1}{E} \left(\sigma_x - \frac{1}{2}(\sigma_y + \sigma_z) \right) \\ \epsilon_y &= \frac{1}{E} \left(\sigma_y - \frac{1}{2}(\sigma_x + \sigma_z) \right) \\ \epsilon_z &= \frac{1}{E} \left(\sigma_z - \frac{1}{2}(\sigma_x + \sigma_y) \right)\end{aligned}\tag{2.2}$$

where σ is the applied stress and ϵ is the resulting strain. Assuming the sensor uses a thin flat rubber sheet with its bottom glued to a rigid mounting surface, the y and z axis of the cube undergo hardly any positional change when a uniform pressure is applied to its surface. That is, the strains ϵ_y and ϵ_z are close to zero. The tactile sensor measures an applied pressure by detecting the positional change in the x direction, which is given by ϵ_x . Using these conditions the resulting x strain can be computed.

First, we obtain the y and z stresses. Since the following relationships approximately hold

$$\begin{aligned}\epsilon_y &= \frac{1}{E} \left(\sigma_y - \frac{1}{2}(\sigma_x + \sigma_z) \right) = 0 \\ \epsilon_z &= \frac{1}{E} \left(\sigma_z - \frac{1}{2}(\sigma_y + \sigma_x) \right) = 0\end{aligned}\tag{2.3}$$

which gives us

$$\begin{aligned}\sigma_y &= \frac{1}{2}(\sigma_x + \sigma_z) \\ \sigma_z &= \frac{1}{2}(\sigma_y + \sigma_x)\end{aligned}\tag{2.4}$$

and finally

$$\sigma_x = \sigma_y = \sigma_z.\tag{2.5}$$

From this we see that

$$\epsilon_x = \frac{1}{E} \left(\sigma_x - \frac{1}{2}(\sigma_x + \sigma_z) \right) = 0.\tag{2.6}$$

This result indicates that the tactile sensor would not detect the applied pressure at all, since no positional change is induced in the z axis. In essence, there is no “escape” path for the rubber to take to relieve the applied pressure. The applied pressure is just converted into internal stresses.

Even if we relax the uniform pressure to the yz surface assumption, using a flat sheet of rubber causes problems. If a point pressure source is applied to its surface, complex positional changes in the z direction will result, making it difficult to infer the actual contact profile. For example, material around the contact area will bulge out of the sensor’s surface.

To overcome this problem, we have formed the elastic-dielectric layer into a sheet with protruding round tabs. As pressure is applied to its top, the tabs compress, and the material expands to fill the surrounding air gaps. Figure 2.1 diagrams this layer. In essence, the tabs permit the material to expand in the y and z directions and allow compression in the X direction.

Now there will be no stress in the yz plane, since it is all transformed into strain. This allows computation of the positional change in the x axis that the previous stimulus will induce:

$$\epsilon_x = \frac{1}{E}(\sigma_x + 0) = \frac{1}{E}\sigma_x.\tag{2.7}$$

This indicates that the tactile sensor equipped with rubber tabs can detect the applied pressure.

Phillips and Johnson [1981] discuss the mechanoreceptors in the human finger. They found that a *plane stress* assumption best fits experimental receptor response data. A plane stress model indicates that all stress is confined to the xy plane of




Figure 2.1: *Elastic-dielectric material. Rubber tabs have been added to the layer to improve its elastomechanical properties.*

the skin. Escape occurs along the z axis. Adding rubber tabs to the tactile sensor makes it more closely approximate the human skin, in this respect.

In addition, the pressure sensitivity of the device is controlled by the properties of the elastic-dielectric material. The stiffness of the rubber, along with the height of the protruding tabs, play a role in the overall sensitivity and the pressure ranges that can be obtained.

To obtain high sensitivity, desirable for detecting slippage and light contact forces, a low stiffness dielectric should be selected. Unfortunately, this reduces the dynamic range of the device. That is, the maximum pressure that the sensor can handle before saturating is reduced. To overcome this, it should be possible to form the tabs with two layers of materials with different stiffnesses, as shown in Figure 2.2. When a force is initially applied to the sensor, the low stiffness portion of the tab compresses. As force is further increased, the higher stiffness material begins to compress. For this scheme to work properly, the low stiffness material must go into non-linear saturation, when it no longer compresses with increasing pressure. The graph shown in Figure 2.2 diagrams its expected behavior.

The dielectric is constructed so each capacitor in the sensor array has a tab




Figure 2.2: *Dual material dielectric layer. By forming the elastic-dielectric layer with two material of varying stiffness, it should be possible to build a sensor with both large dynamic range and high sensitivity for small pressures. The graph shows a plot of pressure against overall tab compression.*

between it. Since the dielectric constant of air is lower than that of rubber, it is best to place the tabs at the capacitor junction sites. The current dielectric has a 0.25 mm thick backing with the tabs protruding out from 0.25 mm to 0.75 mm, depending on the desired force range. The tab diameter is 1.4 mm. Experimental results determined that these values gave suitable performance for an overall device sensitivity in the 0 to 200 grams per square millimeter range.

To form an array of capacitors, electrical traces are attached to the dielectric's upper and lower surfaces, with the upper traces placed perpendicular to the lower ones. A capacitor is formed at each of the trace intersections. A thin layer of electrically insulating silicone rubber is placed around the device to provide electrical shielding. Finally, the sensor is covered with conductive silicone rubber to reduce external electrical interference. Figure 2.3 diagrams a cross section of the device. The next few paragraphs discuss this design in more depth.

The current prototype has 8 upper conductive traces and 8 lower conductive traces, forming 64 force sensitive capacitive cells. The scanning electronics dis-




Figure 2.3: *Tactile sensor cross section. From top to bottom: upper conductive traces, elastic-dielectric material, lower conductive traces. The entire unit is covered with electrical shielding.*

cussed in the following section read force values off the array by detecting each cell's capacitance.

The traces are made of an electrically conductive silicone rubber that bonds to the dielectric layer. The parallel trace pattern is silk screened onto the dielectric layer. The lower dielectric surface must first be made into a flat surface suitable for the traces. To do this, a thin sheet of electrically insulating silicone rubber is bonded to the bottom of the protruding tabs. The lower conductive traces are bonded onto this surface. Currently the traces are 1.27 mm wide, with 0.63 mm gap between adjacent rows. Hence, the capacitor plate area is 1.27 square mm, and the center-to-center cell separation is 1.9 mm.

The sensor must be shielded from stray electrical interference. Without proper shielding, the small capacitances that are to be detected at the trace junctions would be swamped by parasitic affects. To do this, the sensor package is covered with an electrically conductive material. Before this can be done, however, the




Figure 2.4: *Prototype tactile sensor packaging. The dielectric material is mounted onto a printed circuit board. The board provides the lower traces.*

upper and lower traces must be covered with a thin insulating rubber layer, and wires connected to the traces must be brought out from the package.

For initial testing of the device, the upper conductive traces were silk screened to the dielectric layer and mounted onto a printed circuit board (Figure 2.4). The board provides the lower conductive traces. The sensor was then covered with an electrically insulating layer, and finally an electrically conductive layer to provide shielding. Photographs of various parts of the tactile sensor are shown in Figure 2.5.

The materials chosen for the various layers of the sensor were often hard to come by, and had complex interactions when layered together to form the overall device. For example, placing certain silicone rubbers onto each other often interfered with their chemical curing process, preventing one of the layers from hardening. In addition, not all silicone rubbers bond with each other to being with. Iñaki Garabieta, a staff member at the MIT Artificial Intelligence Laboratory, played a crucial role in locating these materials and in constructing the sensor itself. His tremendous




Figure 2.5: Photographs of tactile sensor components. Top from left: printed circuit board, dielectric layer viewed from tab side, dielectric layer and upper conductive traces. Bottom from left: sensor mounted onto printed circuit board with trace connections exposed, back view of sensor and detection electronics.

skill and craftsmanship made it a reality.

2.3.3 Detection Electronics

Before discussing the detection electronics, it is useful to estimate the capacitance of a sensor cell. Hopefully, its magnitude will be large enough to permit detection with relatively simple electronics, keeping the size and cost of the sensor within reasonable limits.




Figure 2.6: Capacitance array schematic diagram.

The capacitance of two parallel plates, in electrostatic units, is given by:

$$C = \frac{\alpha A}{4\pi d}, \quad (2.8)$$

where α is the dielectric constant of the separation material, and d and A are the separation distance and area of the two plates, respectively. If we assume $d = 0.025$ cm, $A = 0.16$ cm², a dielectric constant of 4, and that the farad = 9×10^{11} electrostatic units, the nominal capacitance for a force cell is approximately 0.2 pF. Experimentally, the value was found to be closer to 0.5 pF, possibly due to a higher than expected dielectric constant. This is large enough to detect if ample electrical shielding is provided to reduce external stray capacitance.

Various techniques for measuring capacitance exist. However, not all of them are applicable for this sensor, since constraints exist on the interconnections between the detection electronics and the capacitive cells. This occurs because wires cannot be connected to both terminals of each force sensing capacitor. Thus, before considering how to measure the capacitance of the sensor's array of cells, its effective schematic diagram must be presented.

Recall the basic design of the sensor. Its upper layer contains rows of parallel conductors, placed perpendicular to a lower layer of parallel conductors. The two

layers are separated by the elastic-dielectric material. At each intersection of an upper trace and a lower trace, a capacitor is formed. The overall interconnection pattern results in the schematic diagram shown in Figure 2.6.

This interconnection pattern constrains the schemes available for detecting each cell's capacitance. One can readily see that each small capacitor cannot be connected to a detector alone. In fact, the best that can be done is to select the row line and column line corresponding to the cell. The other cells connected to the selected row and column will affect the output value. The detection electronics must minimize the effects of these other cells. In effect, the detection electronics must isolate each cell in the array from all others, when taking its reading.


Before discussing the actual design chosen for the sensor, a review of several alternative methods may be instructive. Figure 2.7 diagrams several approaches for measuring capacitance, some of which are not applicable to the sensor, but are interesting none the less.

Circuit A diagrams a bridge detection scheme. Here, a time varying signal is applied to the circuit inputs. As the capacitor's value varies, the effective impedance of one side of the bridge changes, affecting the voltage between V_a and V_b . A differential amplifier can be used to detect this difference. The output of this amplifier is proportional to the capacitance being detected. Since a differential signal is being detected and amplified, this scheme has the potential for highly accurate measurements. Unfortunately, it cannot be used with the tactile sensor, since forming the bridge requires access to both terminals of a cell's capacitor.


Circuit B uses the the resulting time constant between the unknown capacitor and a reference resistor. When a unit pulse is applied to the capacitor, the amplifier's output will exponentially decay with a RC time constant. Another circuit, not shown in the schematic, times the decay and then computes the value of C . Since C is quite small, the timing circuit required must be fast and sensitive, complicating the detection scheme's design.

Another approach to measuring changing capacitance is based on charge conservation. To understand how it works, the constituent equation for a capacitor

Circuit A:



Circuit B:



Circuit C:




Figure 2.7: Methods for detecting capacitance.

must be presented. The amount of charge a capacitor can store is given by,

$$Q = CV, \quad (2.9)$$

where Q is charge, C is capacitance, and V is the voltage difference between the capacitor's plates. As the distance between the plates is changed, the charge changes:

$$i = \frac{dQ}{dt} = V \frac{dC}{dt}. \quad (2.10)$$

Here, a current flow of i is induced as the capacitance changes. The voltage is assumed to be constant.

This scheme is realized in Circuit C (Figure 2.7). Here, C_{cell} varies over time and C_{ref} is a fixed reference. Since C_{cell} and C_{ref} are connected, the charge on their plates must be the same:

$$V_{in} \frac{dC_{cell}}{dt} = \frac{dQ}{dt} = -C_{ref} \frac{dV_{out}}{dt}. \quad (2.11)$$

Equation 2.11 holds because the input impedance of an operational amplifier can be made high enough to neglect the current that would flow into its terminal. So, whenever the input capacitor's value changes, a proportional change in the output voltage is induced. Unfortunately, this method does not work well with leaky capacitors, as the ones forming the tactile sensor are likely to be. Over time, charge will actually leak between the capacitor plates, changing the output voltage appreciably.

The detection method chosen for the tactile sensor is also based on circuit C. It's operation is similar to the previous scheme, but overcomes the effects of leaky capacitive cells. Now the input signal V_{in} is varied over time. If we assume C is fixed, and that V varies over time, differentiating Equation 2.9 gives

$$i = \frac{dQ}{dt} = C \frac{dV}{dt}. \quad (2.12)$$

Now, however, the current flow i varies continuously, since the input signal is not constant. Since the junction charge is constantly changing, it is much less sensitive to capacitance leakage than the previous scheme. So, if $V_{in} = \sin\omega t$, the circuit output is given by:

$$V_{out} = -V_{in} \frac{C_{cell}}{C_{ref}}, \quad (2.13)$$

where C_{ref} is the feedback capacitor selected to give an appropriate gain.

Equation 2.13 shows that the output of the amplifier is just a modulated version of the input signal. To measure the capacitance of C_{cell} , the variation in amplitude of the sinusoidal output waveform is detected. First, the output waveform is precision rectified. In essence, its absolute value is obtained. Next, additional gain is applied, and low pass filtering converts the time varying waveform to a direct current signal with magnitude proportional to the root mean square voltage of the original signal. Conveniently, root mean square detector chips can be used to perform most of the detection electronics, greatly reducing the circuit's complexity.

To better understand the sensor's performance, and to select the operating frequency for the input waveform, a more detailed analysis is needed. The frequency of the input signal affects the reactance (generalized resistance) of a capacitor. The reactance of a capacitor is given by:

$$Z = \frac{1}{j\omega C}, \quad (2.14)$$

where ω is the input signal frequency. From this equation we see that the higher the input frequency, the less reactance the capacitor has, and the greater the current flow through it. Since C_{cell} is on the order of 1.0 pf, a relatively high frequency is needed to obtain even a small current. Assuming an operation frequency of 100K hertz, the capacitor has an effective resistance of 10 megaohms. Since the operational amplifier input resistance should be large relative to this, an FET input device and a high frequency signal should be used.

The preceding assumption of infinite operational amplifier gain must be relaxed to fully understand the performance of this detection scheme. Since an amplifier has a finite gain-bandwidth product, high frequency input signals will result in significantly lower amplification. If this is the case, Equation 2.13 will no longer hold.

It is now clear that input signal frequency cannot be increased at will. It will eventually deteriorate the amplifier's performance. On the other hand, the input frequency cannot be too low, or not enough current will flow through the input capacitor. To fully understand this tradeoff, however, we must consider the entire sensor in our model.




Figure 2.8: Capacitance array scanning electronics schematic diagram.

The effective schematic of the capacitive array is shown in Figure 2.6. A scan of the array should measure the capacitance of each of these cells. The scanning electronics must read each cell individually, and must overcome the cross-talk that the interconnect pattern might cause. To do this, an amplifier is connected to each of the sensor's rows, while an input signal is applied to one column at a time (Figure 2.8), in a sequential fashion. Hopefully, the row amplifier outputs will correspond to the capacitance of the selected column's cells.

In Figure 2.9 a more detailed model of the sensor's electronics is presented. Here, the 7 capacitive cells in the column that are not selected are shown connected to the amplifier's negative input, and are labeled C_u . The amplifier's negative input is modeled as a resistor to ground, to indicate that some current flows into it. The resistor R_{cell} models the sensed capacitor's leakage. Applying Kirchhoff's current law to the junction at the amplifier's negative input gives:

$$\frac{\frac{V_{out}}{G} - V_{in}}{Z_{cell}} + \frac{V_{out}}{Z_u} + \frac{V_{out}}{Z_{op}} + \frac{V_{out} - V_{out}}{Z_{ref}} = 0, \quad (2.15)$$

where G is the gain of the amplifier at the particular operating frequency, Z_u is




Figure 2.9: Detailed model of the tactile sensor's scanning electronics.

the effective impedance of the parallel combination of the other capacitors on the selected column, Z_{op} is the impedance of the amplifier's negative input, and Z_{cell} is the impedance of the sensed capacitor and its plate leakage. The impedances are given by:

$$\begin{aligned} Z_{cell} &= R_{cell} / (j\omega R_{cell} C_{cell} + 1) \\ Z_u &= R_u \\ Z_{op} &= R_{op} \\ Z_{ref} &= 1 / j\omega C_{ref}. \end{aligned} \quad (2.16)$$

Solving for V_{out} gives

$$V_{out} = -V_{in} G R_{op} \frac{j\omega C_{cell} R_{cell} + 1}{j\omega R_{cell} R_{op} (C_{ref} G + C_u - C_{ref} + C_{cell}) + R_{op} + R_{cell}}. \quad (2.17)$$

Since FET input amplifiers are used, we can assume R_{op} is very large:

$$V_{out} = -V_{in} G \frac{C_{cell} R_{cell} + 1}{j\omega R_{cell} (C_{ref} G + C_u - C_{ref} + C_{cell}) + 1}. \quad (2.18)$$

Under ideal conditions, where both G and R_{cell} are very large, Equation 2.18 gives the desired result:

$$V_{out} = -V_{in} \frac{C_{cell}}{C_{ref}}. \quad (2.19)$$




Figure 2.10: Graph of the Tactile Sensor's Crosstalk Immunity with respect to amplifier gain. C_u ranges from 0 pf to 20 pf.

The effects of plate leakage can be examined using Equation 2.18. For large gain G , the magnitude of V_{out} is

$$V_{out} = V_{in} \sqrt{\frac{1}{C_{ref}^2 R_{cell}^2 \omega^2} + \frac{C_{cell}^2}{C_{ref}^2}}. \quad (2.20)$$

Errors in the output signal occur when R_{cell} is large or ω is small. Hence choosing a large value for ω minimizes leakage affects.

If we assume the drive frequency is high enough to ensure low plate leakage, Equation 2.18 gives us

$$V_{out} = -V_{in} \frac{C_{cell}G}{C_{ref}G + C_u - C_{ref} + C_{cell}}. \quad (2.21)$$

Ideally, V_{out} should be independent of the capacitors that are not selected (C_u). If we assume that C_{cell} has a capacitance of 1.0 pf, and C_{ref} has a capacitance of 1.0 pf, and V_{in} has amplitude 10, V_{out} should be 10. Any variation from this value will result in an erroneous measurement of the cell's capacitance. Equation 2.21 can be used to determine the gain required for this to be the case.

The graph in Figure 2.10 shows the effects of G and C_u on sensor output, as obtained from Equation 2.21. Sensor output is plotted against increasing gain.

Ideally, the output should be 10, and any deviation from that line is undesirable. Several curves are plotted for different values of C_m , since its actual value depends on the force distribution on the sensor pad. The graph indicates that for gain over approximately 200, adequate performance is obtained. This limits the operating frequency of the input signal to the gain-bandwidth product of the amplifier divided by 200. Using this criterion, a driving signal of 200K hertz was selected for the prototype sensor.

2.3.4 Design Process Summary

The device described in the previous sections required several design iterations before it resulted in a suitable sensor. In fact, the initial prototypes did not work at all. A brief review of some of the earlier attempts and failures at making the capacitive sensor will emphasize the most important features in the final device.

Version one: the lower traces were etched onto a printed circuit board. Conductive paint was used to silk screen the upper traces onto thin rubber materials. Unfortunately, the electrically conductive rubber reacted with the elastic-dielectric layers that it was being applied to, and lost its conductivity. After several attempts, it was decided to apply the conductive paint onto a thin mylar sheet. This worked, but unfortunately mylar is not a flexible material. However, we used this sheet in the initial prototypes in lieu of anything else that worked. The elastic-dielectric materials used were very thin sheets of rubber. The motivation for this was that the junction capacitance would be maximized by a small plate gap. The electronics detection circuitry was fabricated on a wire wrap board, and the sensor was located near the electrical detection components. This device did not work at all.

Version two: evaluation of version one's failure indicated that too little attention had been given to parasitic capacitances. The detection electronics were located far from the capacitive array and the sensor itself was not enclosed within adequate electrical shielding. The next prototype was built with onboard amplifiers, placed in close proximity to the array. Surface mounted components were used to minimize their size. This worked, although the performance was not adequate. Performance problems included large crosstalk between sensor rows and poor response to an

applied force.

Version three: small ground traces were placed between the rows to reduce the cross talk and rubber tabs were added to the elastic-dielectric layer to increase its compressibility. In addition, extensive electrical shielding was added to the entire package. The performance was now satisfactory, approaching the established design goal.

Version four: the mylar sheet with the upper conductive traces had to be eliminated. This layer caused nonlinearities in the sensor's force output and reduced overall accuracy. Since a small, variable amount of air was trapped between the mylar sheet and the elastic-dielectric layer, the actual distance between the capacitor's upper and lower plates was hard to control. After a lengthy search of materials, an electrically conductive silicone elastomer that bonded to the elastic-dielectric layer was found. This resulted in the device described in this chapter.

2.4 Performance Results

A series of experiments was performed to quantify the tactile sensor's performance. The set of experiments selected identify such criteria as pressure sensitivity, shape discrimination, repeatability, stability, hysteresis, and interference immunity. Before beginning to review the experimental results, a brief digression on the sensor evaluation process is presented.

2.4.1 Quantifying a Sensor's Performance

It is easy to compare the performance of two stereo receivers. Over time, people have agreed upon a set of standard metrics that allow contrasting different units. Such figures as "power output" and "harmonic distortion" are commonly used. Without a reasonable and agreed upon convention for these comparisons, it would be difficult to make relative performance statements about different models.

Tactile sensor literature is plagued by a lack of evaluation metrics. Because of this, it is often hard to tell if a new sensor is really an advance over previous efforts. In fact, there is little agreement on any performance evaluation standards at all. For example, no standard metric for reporting a tactile sensor's pressure

sensitivity exists. When pressure is applied to a pad, should it be applied as a point source, or over a large area, or both? To make matters worse, researchers often fail to indicate how a particular test was applied, and report, for example, the sensor's output for a particular force, giving no indication of how that force was distributed over the array.

About the only metric that is consistently used is the number of tactile sensor cells in the array. However, this figure is practically meaningless. Ostensibly, the figure is reported as some indication of the overall spatial resolution of the device. In reality, the spatial resolution is a complicated function of the cell density, the covering materials, and the scanning electronics. It would be useful to have an experimentally determined figure that would overcome this problem, and convey more useful information about the sensor's capabilities.

The psychological literature can provide a useful starting point for quantifying sensor performance. Phillips and Johnson's [1981] have studied human touch sensing extensively using a number of interesting tests. Specifically, they characterize two-point discrimination, gap detection, grating resolution, and letter recognition. While some of their procedures are specific to human subjects, their methodology is of general use for the development of tactile sensor test standards. In addition, use of these tests would allow comparison of tactile sensors to a human's touch capabilities. Perhaps the most important point to be learned from their work is that the design of touch sensing experiments must be done carefully. Since complicated mechanical interactions between the skin, the mechanoreceptors, and the testing apparatus occur, the set of experiments performed must be well thought out for their results to be meaningful.

The following sections present not only a performance evaluation of the sensor, but discuss the significance of the tests themselves. Hopefully they provide a concrete picture of the tactile sensor's capabilities. Direct application of the Phillips and Johnson studies has been omitted due to time considerations.

To gather more accurate and complete data, a tactile sensor tester was developed. The device, diagrammed in Figure 2.11, can apply a force accurately at any point on a tactile sensor pad. A stepper motor controlled *xy* positioning table moves the pad, and a stepper motor controlled linear slide mounted perpendicular to the




Figure 2.11: *Tactile sensor tester.*

table applies the force. The three motors are under complete computer control. If a probe tip were mounted directly on the z axis, a small position changes would produce a large force increase into the sensor since the slide advances a relatively large distances per step. In effect, the granularity of forces that could be applied would not be adequate. To overcome this, a small linear slide is mounted onto the tester's z axis and the probe is attached to the outer side of the slide. A spring is connected between the z axis and the top of the probe. As the probe is pressed into the sensor pad, the spring compresses, and the applied force increases. The linear

slide keeps the position of the probe aligned with the axis. To accurately measure the force being applied, the probe tip is mounted onto a load cell. The load cell's force output is sent to an analog to digital converter for computer processing. The experiments described in the following sections were conducted using this device.

2.4.2 Pressure Sensitivity

One of the most important performance criterion to establish is the sensor's pressure sensitivity. In general, the sensor's response is affected by the applied force *and* the force distribution. Unfortunately, the response per force unit area is not invariant over the applied area. That is, as the area of an applied stimulus of constant pressure is changed, the sensor's output also changes.

The response variation due to stimulus size is due to both mechanical and electrical properties. If the sensor were composed of discrete force sensing elements that have no connection to each other there would be no sizing affect. However, since most sensors have interconnecting components in them (an elastomeric covering, for example), the mechanical connections becomes a factor. The sensor's scanning electronics also play a role in its response. Hopefully the electronics can read a cell without ambiguity. In some sensors, however, the values of neighboring cells can have effects on each other.

The pressure sensitivity measurements obtained do not indicate the precision of the sensor. Instead, they attempt to capture the range of stimuli that the sensor can detect. This is useful for evaluating such factors as how light a touch can the device perceive, or how much pressure forces it into saturation. Since a tactile sensor is not necessarily used as a precise measuring device, this distinction is warranted.

To test the sensor's pressure response, different probe tips with various contact surface areas were pushed into the pad. When a uniform pressure was applied, the sensor response was largely invariant over probe tip size. The rubber tabs in the elastic-dielectric material are responsible for this desirable behavior.

Figure 2.12 show the response of 6 tactile cells subjected to an applied force from 0 to 150 grams. A probe tip with 1.5 mm square area was used. The top

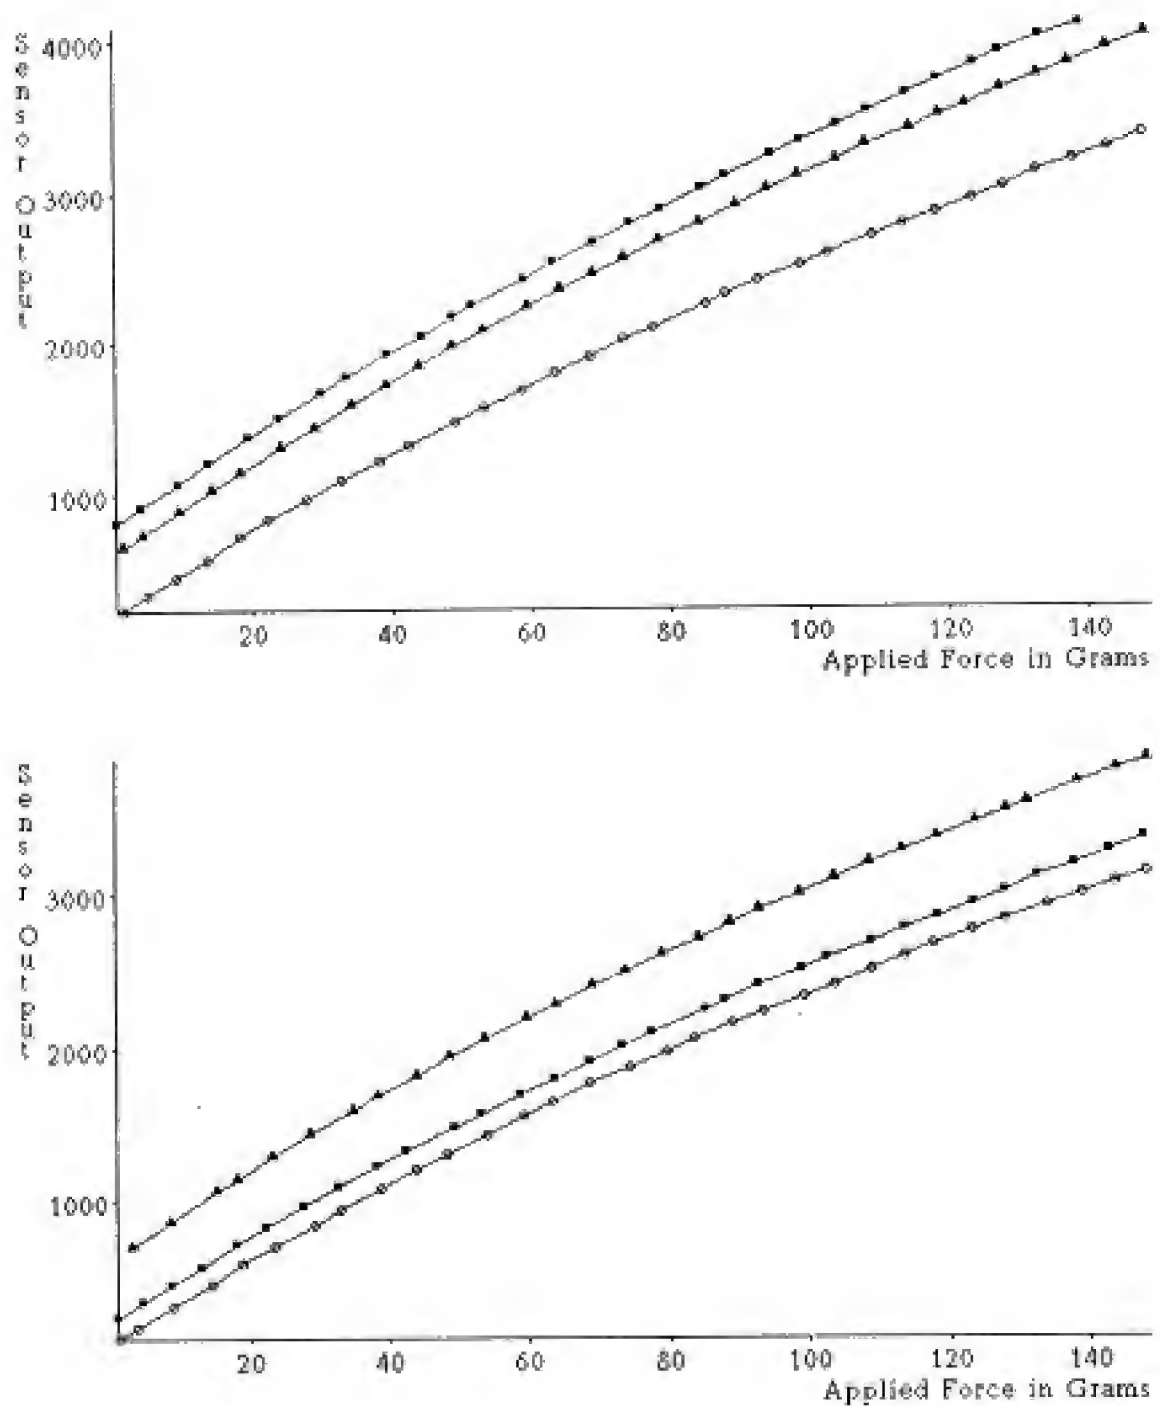


Figure 2.12: Tactile sensor pressure sensitivity. Top: response of 3 cells on same column. Bottom: response of 3 cells on same row.

graph shows the response of 3 cells along the same column of the sensor, while the bottom shows the response of 3 cells along the same row. Each row shares the same detection amplifier and each column shares the same applied input voltage. These plots show that the slope of force response is fairly constant in all cases, though the output offsets vary. Since the variation in cell outputs are similar for both the row case and the column case, no conclusions can be reached on what sources contribute most heavily to the output variation.

2.4.3 *Hysteresis*

Hysteresis plagues most sensors, particularly elastomeric based devices such as the capacitive array. The problem is exhibited by measuring the sensor output as a force probe is pressed and then released from its surface. The output while the probe force is increasing differs from the output while the force is decreasing. In essence, the rubber dielectric expands at a slower rate than it contracts, when subjected to the same force. This is clearly a problem if the sensor is being used to accurately measure an applied force, since it introduces output uncertainty.

In actuality, hysteresis in tactile sensors may not prove to be as major problem as some researchers have suggested. Human skin displays substantial hysteresis, yet our tactile sensing ability is well equipped to handle complex manipulatory tasks. In addition, a robotic system can compensate for hysteresis by recording the time history of the manipulator motion. Since the robot knows if its motion is causing increasing or decreasing force into an object, it can select the appropriate half of the hysteresis curve to translate sensor output into force. Even more importantly, hysteresis is probably unavoidable since tactile sensors must have elastomeric coverings. Previous discussion has established why a soft covering is so important.

To examine the sensor's hysteresis, a probe was pushed into the center of one cell. The sensor's output was sampled while the force was varied from 0 to 150 grams, and then back to 0 grams. Since the speed at which the probe is applied and then retracted affects the level of hysteresis, the force range was covered at several different rates. The plots shown in Figure 2.13 diagram the results of two of the trials. The upper graph corresponds to a force applied over a 2 minute interval

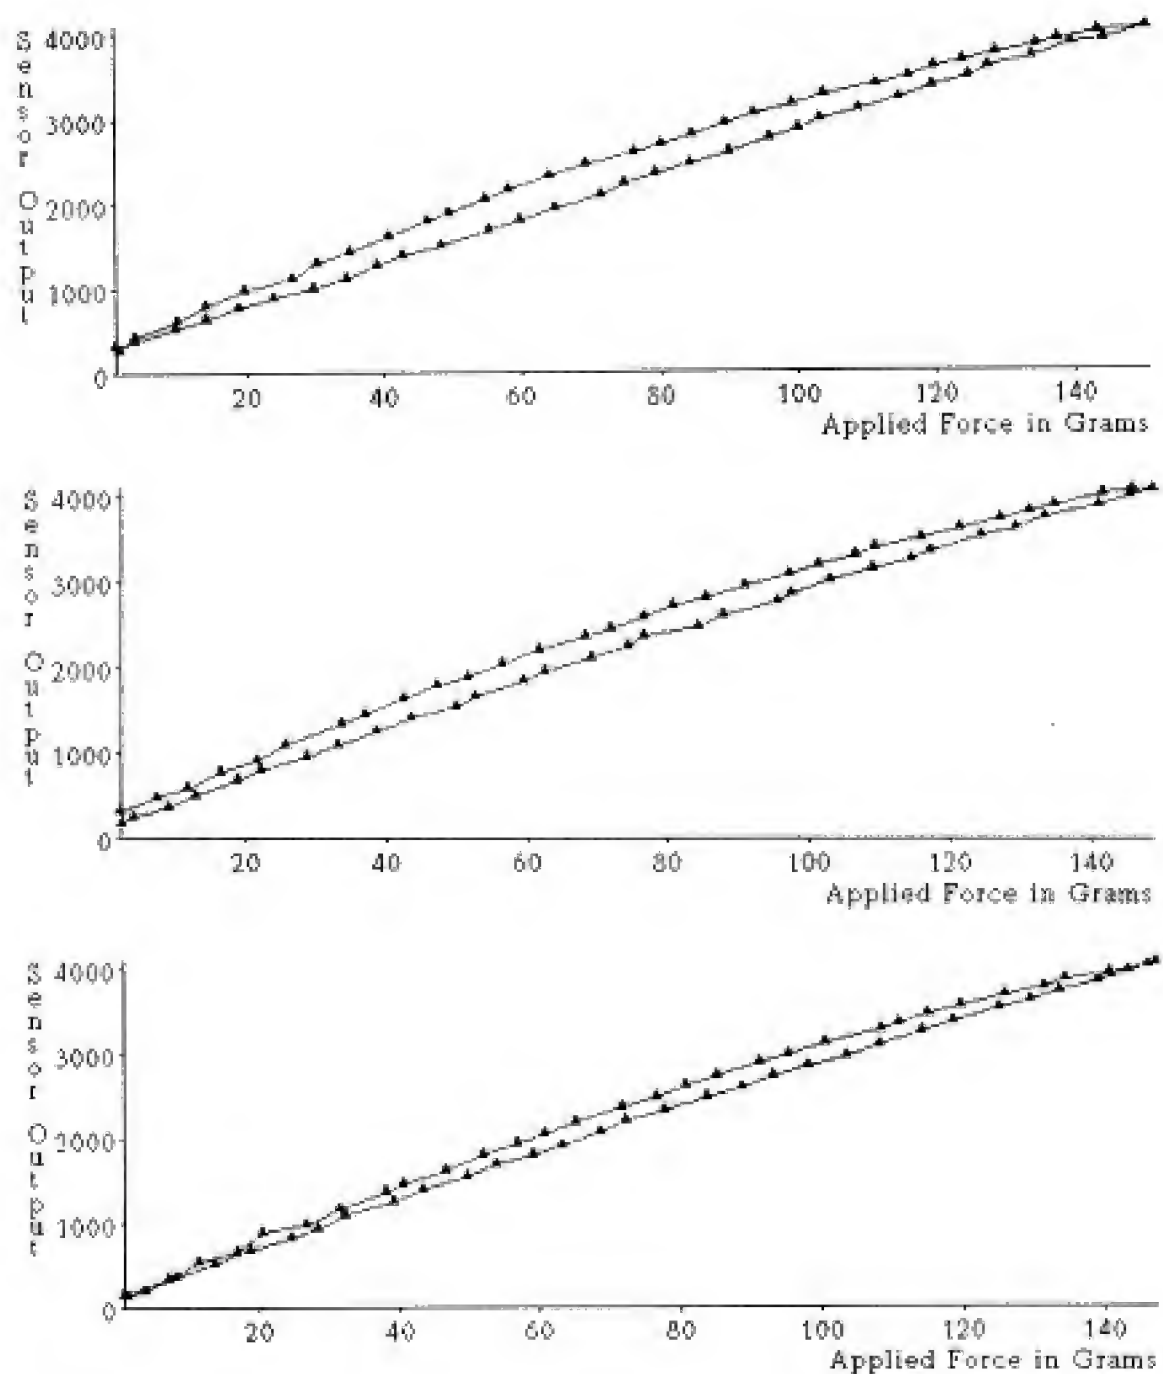


Figure 2.13: Tactile sensor hysteresis. Top: forces applied over 1 minute. Middle: forces applied over 5 minutes. Bottom: forces applied over 10 minutes

while the lower graph was obtained over 5 minutes. As expected, the longer elapsed time reduces the hysteresis.

2.4.4 Measurement Repeatability

Tests were performed to determine how accurately a sensor can measure an applied force. The tester probe was poked into a capacitive cell a number of times, and the sensor's output was recorded. The output variance was used to compute the number of significant bits of force data that the computer was receiving from the pad. Repeated trials found that independent of hysteresis, 8 bits of data were obtained. The actual uncertainty is larger if the force's time history is unknown, since hysteresis is on the order of 5 percent of the full scale reading.

2.4.5 Spatial Selectivity

The spatial resolution of the sensor is a function of the spacing between capacitive cell centers and the elastomeric properties of the dielectric and protective coverings. For a sensor that records just surface normal force, a point source should only be detected by the one or more cells directly in contact with the stimuli; spreading or blurring to adjacent cells should be minimal. Obviously, having a sharp sensor response will give more detail in tactile force outlines of probed objects. In some cases, however, it may be useful to propagate strains between adjacent sensor sites. Fearing and Hollerbach [1984] theoretically show how strain sensors placed below a surface can extract the angle of inclination, location, and magnitude of a load line.

To measure the spatial selectivity of the sensor pad the tactile tester was programmed to step linearly across the sensor pad and to apply a uniform force at each location. The probe was advanced at 0.1 mm intervals and applied 100 grams force at each point. One would expect that a cell would respond best to a force applied directly at its center. As the probe moves further away, a drop in output should be noticed. The receptive field of each cell is affected mostly by the mechanical properties of the rubber elastic-dielectric layer with the tabs helping to localize a cell's response.

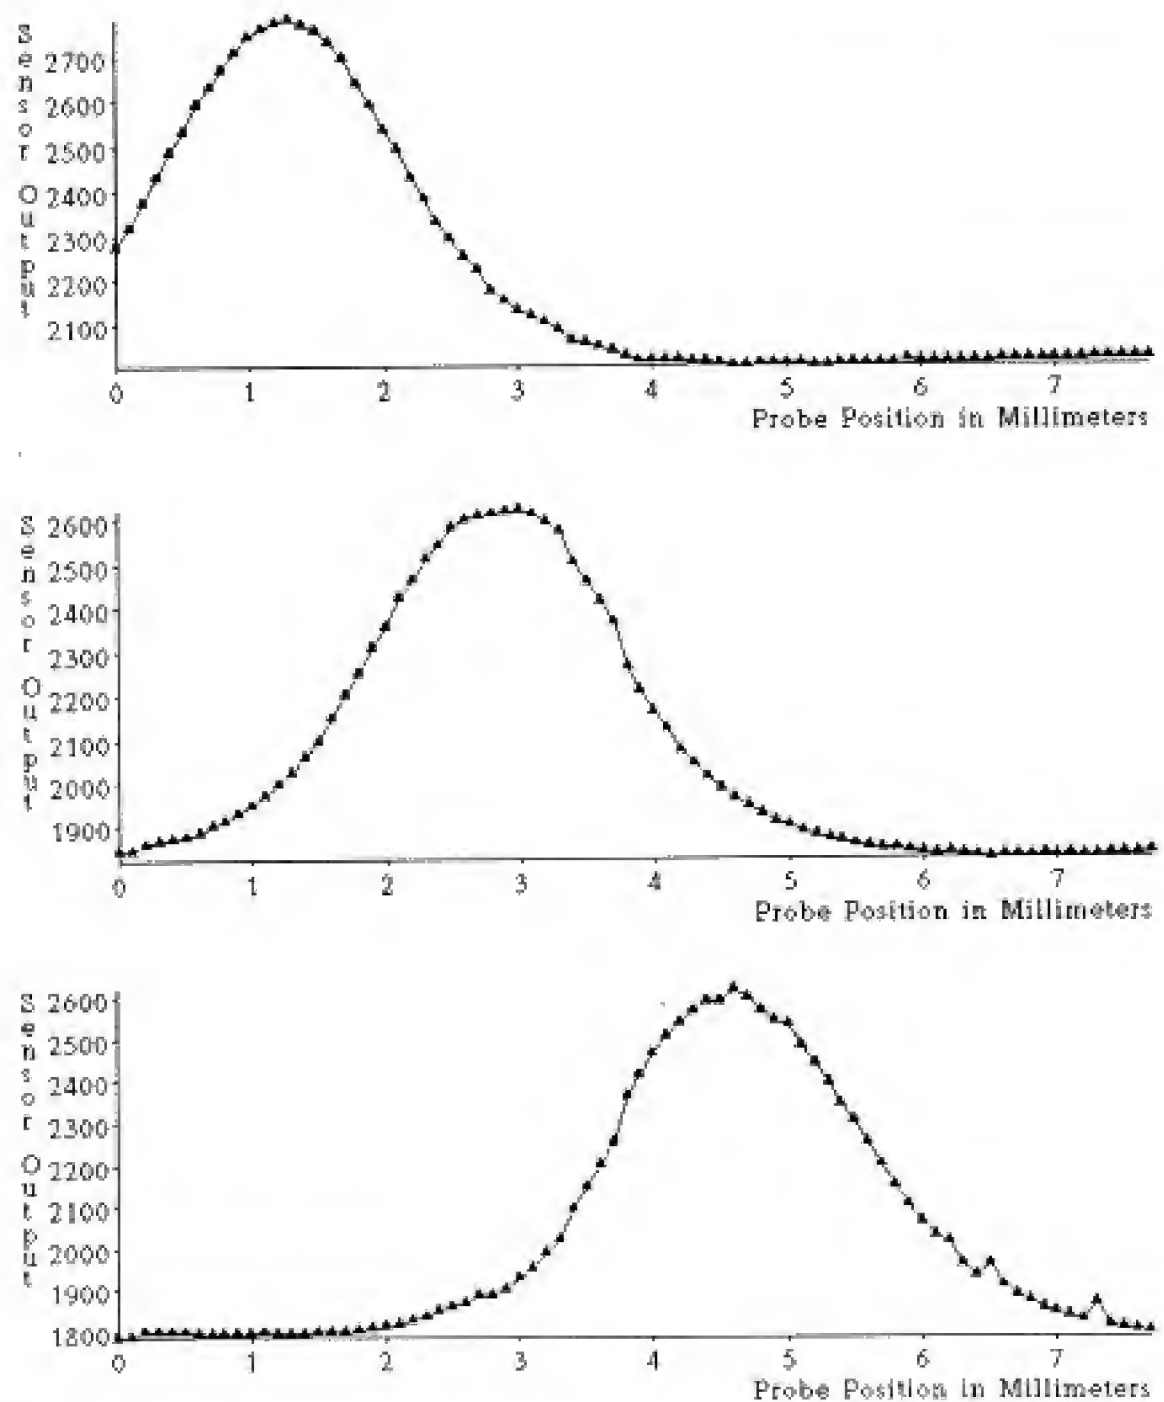


Figure 2.14: Tactile Spatial Selectivity. These curves show the response of sensor cells as a probe is applied to the pad at 0.1 mm linear spaced intervals. The curves, from top to bottom, show the response of the 3 adjacent cells.

The plots shown in Figure 2.14 summarize the results of this test. The curves show the response of 3 adjacent sensor cells as a probe is linearly stepped across them. The peaks of the plots are spaced approximately 2 mm apart as expected, since this is the array's center-to-center spacing. Notice that there is some response overlap between adjacent cells. That is, when probing on a cell, the neighboring cell will show some response as well. This behavior is desirable since it avoids the dead zones between cells that would otherwise result.

2.4.6 *Shape Discrimination*

The shape of an object can be obtained by pressing the tactile sensor against it, and recording the resultant force profile. In cases where visual inspection is impossible, such as when a manipulator end-effector obscures the view, this is especially desirable. The spatial resolution results of the previous section indicate that the sensor's shape discrimination ability should be good. To verify this, small objects were pressed against the tactile pad, and the force outputs were recorded.

Since shape recognition utilizes all 64 tactile cells, a calibration procedure must be employed to normalize the array's response. That is, each cell should read the same value when equal forces are applied to them. Variations in cell output with an applied force occur for several reasons. A cell's force response is related to the compressibility of the elastic-dielectric material at its location. Small variations in the silicone rubber's properties cause differences in compressibility at different pad locations. In addition, the scanning electronics themselves have small component variations, inducing differences in the force output of the cells.

To calibrate the sensor the tactile tester apparatus was employed. The tester was programmed to probe at the center of each tactile cell with a number of different forces. The actual applied force and the cell's output were recorded for each measurement. A least squares procedure was used to fit a straight line to these data points. Table 2.2 shows the calibration data for one of the sensor pads used to record the tactile images shown in this section. The slope and zero force intercept for each cell is given. Since a preceding section showed that the sensor's force response is linear, fitting the data points to a straight line is considered adequate.

Table 2.2: Tactile sensor calibration data. The slope and intercept of the cell output to force output curve. The upper value in each box is the curve's slope. The lower value is the curve's zero force intercept.

11.17 1170.11	9.20 529.50	10.39 984.66	11.43 1059.87	9.44 408.07	9.62 667.39	10.62 861.73	12.62 1359.98
10.46 1315.54	8.66 739.62	10.18 1258.97	10.90 1306.69	8.87 605.33	9.51 849.42	10.51 1111.26	11.51 1616.56
11.52 1556.75	9.76 951.96	11.38 1416.41	11.99 1487.33	9.54 730.60	9.92 991.70	10.91 1146.66	12.83 1708.99
11.84 1526.51	9.95 900.59	10.14 1290.99	11.29 1308.79	9.18 543.93	9.81 854.14	11.05 1081.48	12.55 1637.00
11.87 1500.00	9.78 907.85	10.15 1232.46	11.02 1249.29	9.42 553.22	10.34 873.16	11.23 1146.05	13.33 1650.53
12.28 1649.94	10.08 1013.50	10.85 1402.50	11.41 1340.25	10.30 642.56	10.90 990.80	12.09 1153.65	14.23 1610.83
12.40 1698.85	11.20 956.49	10.90 1270.02	12.20 1241.53	10.20 622.25	10.94 983.01	12.23 1222.92	13.76 1765.72
11.19 1897.31	10.88 946.76	11.69 1304.00	10.64 1139.44	8.53 468.72	10.04 661.06	11.01 836.18	13.01 1347.81

The actual force output is obtained from a sensor's cell output from

$$\frac{c_{out} - b}{m} = f_{out}, \quad (2.22)$$

where m is the curve's slope, b is the zero force intercept, c_{out} is the raw sensor output, and f_{out} is the normalized force. Notice that the values in each column of Table 2.2 are relatively constant. This occurs because each column has its own amplification electronics. Evidently the variations in the amplifier's electronic components contribute heavily to the calibration difference in force output.

Figures 2.15 through 2.23 show tactile force images. The data is displayed using three dimensional bar plots. The height of each of the 8×8 bars corresponds to the cell's force output. Figures 2.15, 2.16, and 2.17 show force images of straight




Figure 2.15: *Tactile images of edges. Top: aligned slightly off horizontal axis. Bottom: aligned directly on vertical axis.*




Figure 2.16: *Tactile images of edges. Top: aligned on a diagonal axis. Bottom: aligned slightly off horizontal axis.*




Figure 2.17: *Tactile images of edges. Top: aligned on diagonal axis, pressed with large force. Bottom: aligned on diagonal axis, pressed with small force.*




Figure 2.18: *Tactile image of integrated circuit socket. Top: diagram of socket. Bottom: image of socket.*




Figure 2.19: *Tactile images of integrated circuit socket. Top: aligned on diagonal axis. Bottom: aligned on opposite diagonal axis.*




Figure 2.20: Tactile image of part with two edges. Top: diagram of part. Bottom: side view image of part.




Figure 2.21: *Tactile images of part with two edges. Top: aligned on vertical axis. Bottom: shifted right.*




Figure 2.22: Tactile images of bolt and washer. Top: bolt. Bottom: washer.




Figure 2.23: *Tactile image of a Molex connector. Top: diagram of connector. Bottom: tactile image.*

edges pushed into the sensor pad at different angles. Figures 2.18 and 2.19 show images of an integrated circuit socket. Figure 2.20 and 2.21 show images of a fairly complex part with two parallel edges, one thicker than the other. The actual gap between the edges is approximately 3 mm, corresponding to the 1 pixel gap shown in the image (sensor cells are spaced 1.9 mm center-to-center). Figure 2.22 show images of a bolt and a washer. Finally, Figure 2.23 shows the profile of a 6 pin molex connector. In all cases the variations in pixel height along an object's edges come from variations in the pressure distribution as the object was pushed into the sensor pad. No special provisions were taken to insure that a uniform pressure was being applied over the part.

2.4.7 Interference Immunity

It is important to insure that the tactile sensor will function in a variety of working environments. One problem with capacitive sensing frequently cited in the literature is their sensitivity to external interference. In essence, parasitic capacitance can interfere with the sensor's operation. To overcome this problem, the device constructed is extensively shielded; the entire sensor is enclosed within a grounded chamber.

To test the array's response to external stimuli, a series of objects were placed near it, without making direct contact to its surface. The sensor output was monitored for any detectable variations due to these external perturbations. None could be found. The device's shielding adequately protects it from external interference.

2.5 Overall Device Evaluation

The beginning of this section discussed the lack of a set of common benchmarks for sensor comparison. This, of course, makes it hard for the results in this section to be compared with other works. However, a general discussion of how this tactile sensor compares with previous works is warranted.

The sensor has achieved significant advances over Boie's prototype tactile sensor. The cell center-to-center spacing has been reduced from 2.5 mm to 1.9 mm, a 32 percent reduction. The technology employed within the sensor should support a

further reduction of cell spacing to the level of 1 mm. Of equal importance are the improvements to the sensor's packaging. The fully shielded integral silicone rubber package gives the device increased sensitivity and reliability, and permit mounting on curved surfaces. The use of tabs in the silicone rubber elastic-dielectric layer give much more flexibility in selecting pressure sensitivity than the nylon stocking used in Boie's sensor. The scanning electronics have been redesigned to improve their performance and reduce their size, to facilitate mounting the sensor in smaller spaces such as on the fingers of the Utah-MIT hand. Finally, the sensor can be scanned faster than Boie's, at a rate of 500 frames per second.

The primary design goal of making a tactile sensor suitable for use with the Utah-MIT hand has been achieved. The sensor can be mounted on the curved finger surfaces, and the size of the electronics is within the space available at the fingers. The cell spacing of the current prototype is close to the 1.5 mm center-to-center spacing that was considered desirable.

Unfortunately, the sensor only responds to surface normal force, and not tangential force components. This will complicate certain tasks, such as slip detection. The sensor's hysteresis might also be a problem since it reduces the overall sensitivity of the device. The hysteresis problem may be overcome by employing a different elastic-dielectric material. Shear force detection, however, does not seem possible using this sensing technology.

The sensor's performance approaches some of Harmon's goals. Its cell spacing approaches the suggested 1.5 mm range, its hysteresis is near the suggested 10 percent level, and its force sensitivity exceeds the 40 grams per millimeter recommendation.

Future improvements planned for the device include a reduction in cell size to the level of 1.2 mm center-to-center spacing, increasing the array size to 16×16 cells, and fabricating the detection electronics using hybrid circuit construction technology. All of these improvements should be possible without any major advances over the current prototypes.

Thermal Sensing

A visual observation of a smooth sheet of red metal, a smooth sheet of red wood, and a smooth sheet of red plastic would not uncover any major differences between them. After touching them, however, it becomes clear that they are not the same. The materials will feel different even if their surface textures are identical. It is the difference in their thermal properties that allows us to tell them apart.

The human finger's sensing abilities include more than just tactile reception. Clearly, metal feels cooler than wood and plastic. This sensation has nothing to do with the absolute temperatures of the materials; if they are in the same room they are indeed likely to be at the same temperature. The sensation is in response to the heat conduction properties of metal compared with wood and plastic. The finger is warm; a constant supply of blood acts as a heating source. When contact is made with a good thermal conductor such as a metal, heat quickly flows out of the finger. This reduces the temperature of the finger, and hence the metal feels cool.

This phenomenon is explained by the second law of thermodynamics. When two materials are placed in contact with each other heat will flow from the warmer one to the cooler one until their temperatures equalize. This heat flow can occur

by three unique forms of physical heat transfer. Heat flow through connected parts of a body is called *conduction*, heat flow through unconnected parts of bodies is called *radiation*, and heat transferred by the relative motion of bodies is called *convection*.

The flow of heat in a thermal system is analogous to the flow of current in an electrical system, where a voltage difference induces a current flow. Voltage continues to flow between two points until the potential difference is eliminated.

The robotic thermal sensor is patterned after a human finger's capabilities. The device should be able to distinguish materials by their thermal conduction properties alone, by measuring the rate at which a material can absorb heat from the sensor. By equipping a robot with an artificial thermal sensor, material classification that visual and tactile sensing cannot accomplish, will be possible. In the next section potential robotic uses for this device are discussed in more detail.

3.1 Robotics Applications for Thermal Sensing

An intelligent robotic system can exploit thermal sensing capabilities in a number of ways. For example, the material in contact with the sensor can be identified by measuring its thermal conductivity properties. If the material type is already known, the sensor output could indicate the roughness of the material's surface texture; a smooth surface would form better contact with the sensor's surface, giving better heat conduction between them.

A thermal sensor can also be used for extracting tactile like force response. To some extent, as the contact force between a material and the sensor increases, the rate of thermal conduction between them will also increase. The actual materials used in the sensor's construction affect the magnitude of this response. If the sensor is ultimately planned for material identification, this pressure response should actually be minimized. These issues will be covered in more detail in the following section.

A two dimensional planar array of thermal sensing elements will give heat conduction images of contact regions. With this, the point of contact between a material and the sensor can be determined. In addition, objects composed of more

than one type of material can be more readily recognized. The orientations of a writing implement composed of more than one material could be determined by probing at several different locations and recording the types of materials found at each point.

Grimson and Lozano-Perez [1985] have shown that very few sensor data points from an object in three space are required for its unique identification. Their work concentrates on the constraints that tactile data provide. However, it would be easy to extend the methodology to include other types of sensor information, including those obtained from thermal sensors. For example, if an object is grasped by a robotic hand equipped with tactile sensors, additional thermal sensors can further reduce the possible object orientations that are consistent with the overall sensor data.

Equipping a robot with both tactile and thermal capabilities will allow a more detailed study of human contact sensing abilities. The interaction between tactile and thermal transduction in our skin is not very well understood, and it is likely that a complex interaction between the two gives rise to our perception of surface textures. To better understand how humans perceive and use contact sensing information, possible theories that attempt to mimic some of the biological operations that the human system performs can be developed and tested on robotic devices.

When this project was first undertaken, a literature search revealed only one reference to a thermal sensor [Dario et al. 1984]. Dario's sensor is based on the pyroelectric properties of polyvinylidene fluoride. Though his device is interesting, it is designed primarily for use as a tactile sensor. At the completion of this work a concurrent project conducted by Russell [1985] was also reported in the literature. It seems that our independent efforts confirm the same result, that thermal sensors can provide useful feedback information.

3.2 Design of the Thermal Sensor

To measure heat conduction, the thermal sensor combines an active heat source with an array of temperature sensors. The temperature transducers are mounted at the surface of the sensor, and record the temperature at the junction between the

Table 3.1: Thermal properties of materials.

Material	K (Btu/hr ft F)	c (Btu/lb F)	ρ (lb/cu ft)	α (sq ft/hr)
Air	0.014	0.24	0.81	0.72
Aluminum	117	0.208	169	3.33
Asbestos	0.087	0.25	36	0.01
Brass	56	0.92	532	1.14
Copper	224	0.91	558	4.42
Lead	20	0.03	705	0.95
Nickel	34.5	0.103	555	0.60
Rubber	0.087	0.48	75	0.002
Silver	242	0.56	655	6.6
Wood	0.12	0.57	51	0.004

sensor and the material being sensed. If contact with a good thermal conductor is made, heat will rapidly flow from the sensor into the material, reducing the temperature at the sensor's surface. The rate of reduction in sensed temperature, and the steady state temperature that the sensor assumes, is related to the thermal conduction properties of that external object.

The following sections explore this basic transduction principle in more detail. First, a theoretical examination of the sensors performance abilities is presented. Next, a discussion of the design and construction principles employed follows. Finally, an experimental performance analysis is conducted, with the results contrasted against the theoretical model of the sensor.

3.2.1 Theoretical Analysis of Sensor Performance

To better understand the performance of the thermal sensor, a theoretical analysis of its behavior is presented. A simplified model of the sensor and the material being sensed that captures the essential parameters that affect the system's response is used. As we will later see, the actual output of the sensor closely corresponds to the results obtained in this section.

To formally classify a material's thermal properties, the parameters of conductance and diffusivity are defined. The coefficient of conductance, K , relates a temperature differential to a heat flow, according to

$$-K \frac{\partial v}{\partial x}, \quad (3.1)$$

where v is the temperature and x is the position along the material's axis. This quantity captures the intuitive notion that some materials conduct heat better than others; the higher their conductance, the lower the temperature difference required for a particular heat flow.

A material's diffusivity is defined to be

$$\alpha = \frac{K}{\rho c}, \quad (3.2)$$

where K is its conductance, ρ is its density, and c is its specific heat. This parameter is related to the change in temperature of a unit volume as a unit temperature difference allows heat to flow into the object over a unit time interval. The units of α , square feet per hour, confirm this. Table 3.1 lists the thermal characteristics of a number of common materials. The wide variety of values of conductivity and diffusivity for common substances should ensure different sensor responses for different materials, and hence allow unique identification.

The thermal sensor is modeled as a block, and the material being sensed as an infinite rod, as shown in Figure 3.1. We are interested in finding the temperature at the boundary between the sensor and the material being sensed. This will model the experimental response of the thermistors in the actual sensor. The material in this section is developed from the excellent discussion of heat conduction problems in Carslaw [1948].

The temperature, v , of a block of material in one dimension, is given by the diffusion equation

$$\frac{\partial^2 v}{\partial x^2} - \frac{1}{\alpha} \frac{\partial v}{\partial t} = 0, \quad (3.3)$$

where α is the thermal diffusivity of the material. To reduce equation 3.3 from a partial differential equation to an ordinary differential equation, we take its Laplace transform

$$\int_0^\infty e^{-st} \frac{\partial^2 v}{\partial x^2} dt - \frac{1}{\alpha} \int_0^\infty e^{-st} \frac{\partial v}{\partial t} dt = 0, \quad (3.4)$$




Figure 3.1: Thermal sensor model. The block on the left represents the sensor. The block on the right represents the material being sensed.

which becomes

$$\frac{d^2 \mathbf{v}}{dx^2} - \frac{s}{\alpha} \mathbf{v} = -\frac{1}{\alpha} V_0(x), \quad (3.5)$$

where V_0 is the initial temperature distribution on the bar, which we assume to be zero.

The boundary condition for equation 3.3, which describes the interface between the block and the rod, is given by

$$Mc \frac{du}{dt} - K \frac{\partial v}{\partial x} = Q, \quad (3.6)$$

where $u(t)$ is the temperature of the sensor block, M is the mass of the sensor, c is the specific heat of the sensor, K is the thermal conductivity of the material the sensor is in contact with, and Q is the heat the sensor is producing. Intuitively, this equation indicates that some of the heat produced by the sensor is absorbed by the sensor itself, and the rest flows into the material being sensed. Since the rod is assumed to be infinite in length, no boundary condition for the other end is needed. We can assume that at time $t > 0$, since the sensor makes good contact with the material,

$$u(t) = v(0, t), \quad (3.7)$$

which allows us to obtain the Laplace transform of the boundary condition in terms

of v ,

$$Mcsv(s) - K \frac{dv}{dx} = \frac{Q}{s} - McV_0, \quad \text{at } x = 0. \quad (3.8)$$

We now must find a solution for $v(s)$ which satisfies boundary condition 3.6 and the governing equation 3.4. A solution that works is

$$v(s) = \frac{Qe^{-rs}}{Mc\alpha s(r+h)} + \frac{V_0 e^{-rs}}{\alpha s(r+h)}, \quad (3.9)$$

where

$$h = \frac{K}{Mc\alpha},$$

and

$$r^2 = \frac{s}{\alpha}.$$

Taking the inverse Laplace transform of equation 3.9, we can obtain $v(x, t)$:

$$\begin{aligned} v(x, t) = \frac{2Q}{K} \sqrt{\frac{\alpha t}{\pi}} e^{-\frac{x^2}{4\alpha t}} - \frac{Q(1+h x)}{Kh} \operatorname{erfc}\left(\frac{x}{2\sqrt{\alpha t}}\right) + \\ \frac{Q}{Kh} e^{hx+h^2\alpha t} \operatorname{erfc}\left(\frac{x}{2\sqrt{\alpha t}} + h\sqrt{\alpha t}\right) + V_0 e^{hx+h^2\alpha t} \operatorname{erfc}\left(\frac{x}{2\sqrt{\alpha t}} + h\sqrt{\alpha t}\right), \end{aligned} \quad (3.10)$$

where erfc is the Gaussian error function, which is defined by

$$\operatorname{erfc}(x) = 1 - \frac{2}{\sqrt{\pi}} \int_0^x e^{-\phi^2} d\phi. \quad (3.11)$$

At the point $x = 0$, the above equation represents the temperature variation of the thermal sensor itself,

$$u(t) = \frac{2Q}{K} \sqrt{\frac{\alpha t}{\pi}} - \frac{Q}{Kh} + \frac{Q}{Kh} e^{h^2\alpha t} \operatorname{erfc}(h\sqrt{\alpha t}) + V_0 e^{h^2\alpha t} \operatorname{erfc}(h\sqrt{\alpha t}). \quad (3.12)$$

From equation 3.12 we can obtain the general behavior of the thermal sensor when it is placed in contact with a material. Since the erfc function approaches zero rapidly, the initial value of the sensor when placed in contact with an object will fall from V_0 to

$$\frac{2Q}{K} \sqrt{\frac{\alpha t}{\pi}} - \frac{Q}{Kh}. \quad (3.13)$$

Figure 3.2 shows the behavior of equation 3.12 as α and M are varied.





Figure 3.2: Behavior of the thermal sensor's governing equation.

3.2.2 *Performance Issues*

The design of a thermal sensor is related to the thermal properties one is most interested in sensing. This chapter will cover sensors that use heat conduction measurements for position localization, and material identification.

Several factors affect the sensor's accuracy when used for material identification. Ideally, the sensed object should be in firm contact with the sensor. The heat flow from the sensor to the sensed object should be entirely unimpeded. This means that increased pressure of the sensed object to the sensor should have a negligible effect on the rate of heat conduction between them. In essence, the thermal resistance of the gap between the sensor and the sensed object should be minimal.

Selection of an appropriate covering material for the sensor can assist in achieving good thermal contact between the sensor and an external object. A slightly compressible material that conforms to the surface of an object seems ideal. If the material deflects substantially with pressure, however, the thermal sensor will have a significant pressure response. While this would be undesirable for material identification, it does allow construction of a sensor with tactile response. The more compressed the sensor covering becomes, the better the thermal contact between the external object and the sensor's transducers. Hence, the sensor gives a tactile pressure response.

The covering material must also provide protection for the sensor's transducers. The electronic components that measure the temperature at the surface of the sensor should be shielded from the damage that pressing the sensor against objects could cause. In addition, the covering must provide a surface that is adequate for the overall functions of the end-effector that it is mounted on; while the sensor might be a device just used to gather exploratory information, it is more likely to be a part of an overall robotics system, and incorporated into existing end-effector designs. It will need to function not only as a sensor, but as the gripping surface for the robotic device.

When selecting the covering material for a multipoint array sensor other issues must be considered. A material with high thermal conductivity will have the desirable property of conducting heat in an unimpeded fashion from the sensor

to the external object. However, it may also be a good heat conductor between adjacent sensor pixels. This will, in effect, blur the readings of nearby points, and reduce the overall spatial sensitivity of the sensor.

To overcome this blurring effect, it is desirable to select covering materials that conduct heat best in the vertical direction. Certain materials that we are looking at have properties that meet this requirement, including a silicon elastomer embedded with lateral conductive strips [Chomerics]. Selecting very thin materials that are relatively poor thermal conductors may also help. While these material reduce the rate of heat flow between the sensor and the sensed object, and hence reduce the overall sensitivity of the sensor, they minimize blurring. A very thin material's lateral insulation effect would be minimal.

The actual medium that the sensor's temperature transducers are mounted on, and the location of the active heat source, are two additional important design factors. The sensor transducers are mounted on some material which, of course, conducts heat between adjacent pixels. In addition, the mounting material will act as a thermal reservoir, reducing the overall response speed of the sensor.

When an object comes in contact with the sensor, heat is conducted off both the temperature transducers and the sensor's overall packaging. The more heat the entire sensor stores, the slower its response rate. Ideally, if the transducers were entirely isolated, and they supplied their own source of heat, contact with an object would draw heat directly from them, and would be immediately measured.

This section has provided information useful for constructing a thermal sensor. The following sections discuss the scanning electronics and the fabrication process employed for the prototype device.

3.2.3 Fabrication Details

The thermal sensor uses thermistors for its temperature measurement transducer. Thermistors are semiconductor devices whose resistance changes with temperature according to the equation

$$R_t = R_0 + K_{therm} \delta t, \quad (3.14)$$




Figure 3.3: Thermal sensor printed circuit board layout. The upper layer of the board is shown in solid lines. Each thermistor is connected to an upper and lower trace.

where R_t is the resistance at temperature t , R_0 is the thermistors nominal resistance, and K_{therm} is the coefficient of resistance with temperature. Typically, K_{therm} is approximately $-0.04/^\circ\text{C}$, giving a reasonable variation in resistance for small temperature changes [Horowitz 1980].

Thermistors are appropriate for several reasons. Most importantly, they are available in small surface mount packages, allowing placement at close intervals. In addition, thermistors have a relatively large variation of resistance with temperature, greatly reducing the complexity of the detection electronic, and reducing the device's overall size. Finally, their linear response curve simplifies the calibration procedures required.

A 4×4 array of thermistors, providing 16 temperature sensing elements, is attached to a flexible printed circuit board. The thermistors are placed 3.5 mm apart, giving the sensor a 10.5 square mm area. Use of a flexible printed circuit board allows the sensor to be mounted on curved surfaces such as the fingers of the Utah-MIT hand. Each thermistor must be connected to a row conductor and a column conductor to allow the cell to be scanned by the detection electronics. The top surface of the circuit board provides the upper conductors, while the bottom




Figure 3.4: *Photographs of the thermal sensor. Left: exposed view of the thermistors attached to the flexible printed circuit board. Right: view after thermally conductive covering has been applied.*

surface provides the perpendicular lower traces (see Figure 3.3). The thermistors are attached to the circuit board using either solder or a conductive glue. The glue is somewhat easier to apply to the small devices, and provides an adequate bond.

A photograph of the exposed thermistors attached to the flexible printed circuit board is shown in Figure 3.4. To smooth the surface of the device and fill the gaps between the discrete thermistors, a layer of thermally conductive silicone rubber is formed over the entire surface of the board. This not only protects the thermistors, but it averages out the temperature response between the gaps present in the thermistor array. The lower photograph in Figure 3.4 shows the sensor after the conductive rubber has been applied. The electronic components seen in the picture are for the detection electronics of the tactile array which is to be mounted above the thermal sensor in the near future.

The thermal sensor's heat generating source is placed behind the flexible printed circuit board. The heater uses an electrically conductive paint to provide resistive heat. Wires are attached to opposite sides of a layer of conductive paint, giving




Figure 3.5: *Thermal sensor cross section. From top to bottom: thermally protective covering, thermistors, flexible printed circuit board, electrical shielding, heating pad.*

a uniform current flow, and hence providing uniform warmth. The heater cannot be mounted directly onto the back of the circuit board, since it would short out the conductive traces. Instead, a thin layer of electrically insulating but thermally conductive paint is applied between their two surfaces.

To summarize this section, Figure 3.5 shows a cross section of the sensor. The top layer of thermally conductive rubber forms the sensor's outer covering. Next come the thermistors, attached to the underlying flexible printed circuit board. The back of the board is covered with an electrically insulating paint. Finally, a resistive paint is applied over this covering to provide the heating source.

3.2.4 Detection Electronics

To minimize the number of interconnection wires that are needed to drive the thermal array, and to detect the resistance of each of the thermistors, a matrix

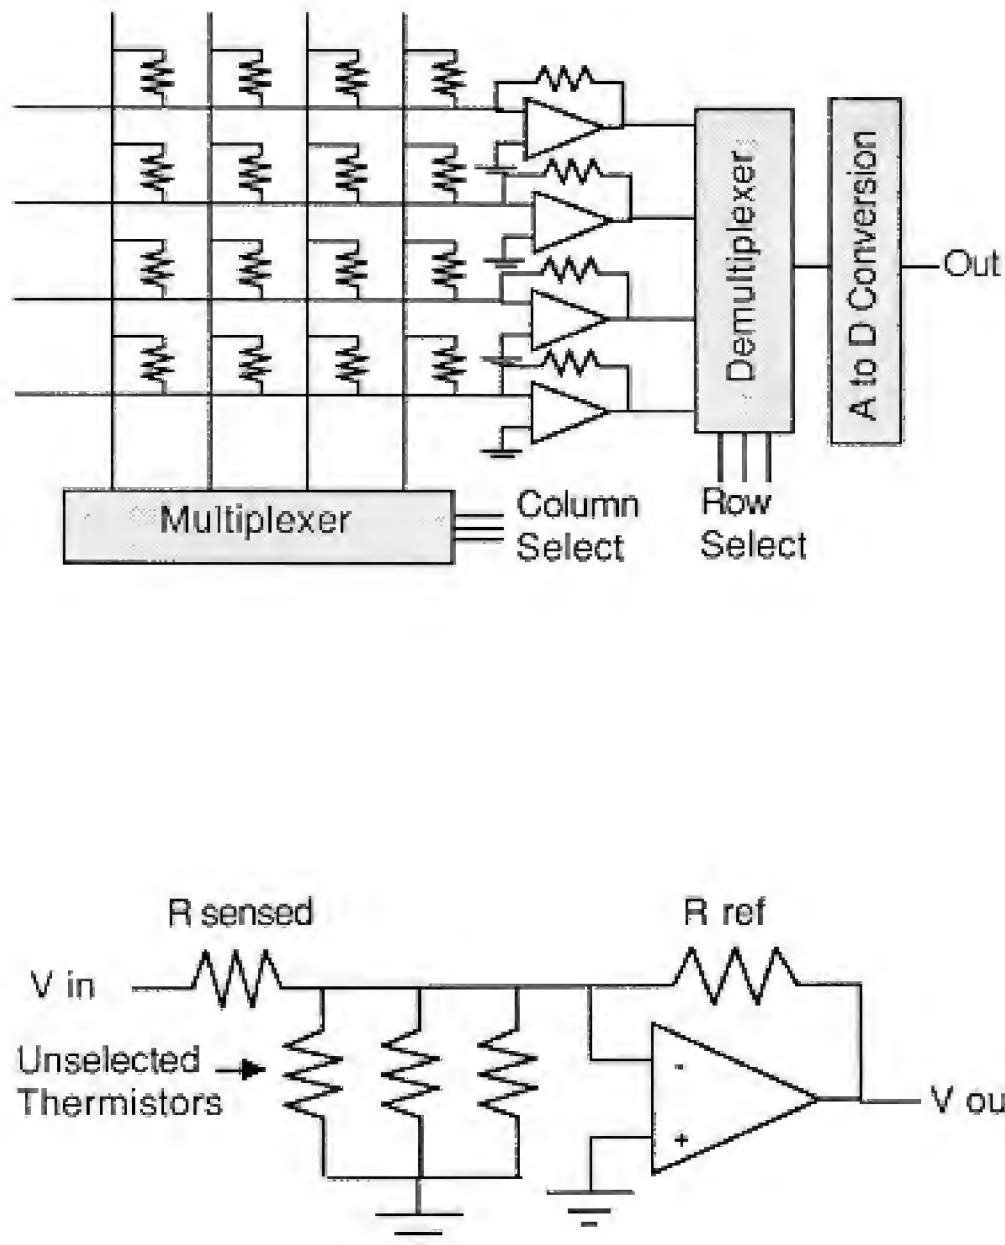


Figure 3.6: Thermal array scanning electronics. Top: overall schematic of the thermistor array and the multiplexing and detection electronics. Bottom: effective schematic when one column has been selected.

scanning approach is employed. Only one wire for each row and each column of the sensor is needed. For the 4×4 sensor, this reduces the number of wires required from 32 to 8, and is obviously desirable. The overall electrical schematic for a 4×4 array is shown in the upper diagram in Figure 3.6. Notice that one terminal of each thermistor is connected by parallel wires and that the other terminal is connected by a perpendicular set of parallel wires. This interconnect scheme can easily be realized with a two sided circuit board.

It is possible to isolate the resistance of any of the thermally sensitive resistors in the array. To read the value of a particular transducer, the column that it is located on is selected with the column select demultiplexer. This will apply a fixed voltage to that column, and tie the other columns to ground. The rows of the sensor are connected to amplifiers in an inverting amplification configuration. Since the positive input of each of the amplifiers is tied to ground, the negative input will be, in effect, at a virtual ground. This allows the resistance of the selected thermistor to be read without ambiguity.

Figure 3.6 shows the effective schematic when a particular column of the sensor is selected. From this we see that only that particular column of thermistors can have an appreciable effect on the row amplifier's output. Since the negative input of the amplifier is at a virtual ground, each of the unselected thermistors are effectively tied to ground and no appreciable current can flow through them. The resistance of the thermistor is obtained from the output of the op amp, and is given by the well known equation,

$$V_{out} = -V_{in} \frac{R_{ref}}{R_{sensed}}, \quad (3.15)$$

where R_{sensed} is the resistance of the selected thermistor, and R_{ref} is the op amp feedback resistor.

In selection of the nominal resistance values for the thermistor it is important to consider their self heating effects. The thermistor will generate a detectable amount of heat flow if its nominal resistance is too low, which can interfere with the sensing process. Hence, the nominal resistance of the thermistors should be high. For our prototype, 5000 ohms was chosen.

By self heating the thermistors it is possible to eliminate the need for a separate

heating layer. If the resistance of the thermistors are low, the current flow through them will generate heat. The scanning electronics could be modified to maintain a constant voltage across each thermistor by regulating the current flow. The current flow itself would be used to determine the heat conductivity between the sensor and the sensed material. Unfortunately, small thermistors with a low enough nominal resistance to generate an appropriate level of self heating were not available. Since the performance of the sensor would be severely diminished by reducing its heat output, a separate heating layer is required.

3.3 Experimental Result

To evaluate the performance of the prototype sensor a series of experiments were conducted using the previously described 4×4 thermistor array. The tests investigate the sensor's sensitivity, repeatability, response speed, and spatial accuracy. For each of the experiments, a graphical representations of the sensor output plotted against time is shown.

3.3.1 Material Identification

One of the primary uses for the thermal sensor is recognition of an unknown material from a library of thermal profiles. From the previous discussion we can expect that the temperature at the surface of the sensor will drop at an exponential rate when placed in contact with an object. The shape and final value of this curve, as we now know, is related to both the diffusivity and conductivity of the sensed object, in addition to characteristics of the sensor itself. Material recognition can be established by matching the sensor's response to a library of response curves, where a good match constitutes identification.

To measure the sensor's effectiveness in distinguishing among different objects, the following experiment was conducted: a number of objects were placed in contact with the sensor, and the temperature over time response was recorded. Figure 3.7 shows plots of the response of one sensor pixel after contact with the materials have been made. The sensors temperature response to wood, aluminum, steel, and brass is shown to differ significantly enough to distinguish among them. Though




Figure 3.7: Thermal sensor response to various materials. Top: response for wood, aluminum, and steel. Bottom: response for aluminum and brass.

the elapsed time shown on the plot is fairly long, positive recognition can be accomplished more rapidly. The temperature curves separate from each other after only a few seconds.

Figure 3.8 shows the temperature response profile for aluminum over a long time period. It takes over 200 seconds for the sensor to reach its steady state condition. Notice that after 250 seconds, the sensor's temperature starts to increase. This behavior was predicted in the theoretical model developed previously. It should be possible to obtain an unknown material's thermal parameters by fitting




Figure 3.8: *Thermal steady state response curve for aluminum. Over 200 seconds are required for aluminum to reach its steady state temperature. Notice that the sensor's temperature eventually starts to increase.*

a theoretical response curve to the sensor's output.

It should be noted that the surface texture of a material being sensed affects its overall thermal profile. This interesting result is caused by variations in the quality of the thermal contact made between the sensor and the sensed object. A rough texture has many gaps in its surface. The gaps are filled with air, and form a thermal insulation layer between the sensor and the object. Since material identification is made by comparing a thermal response curve with a library of such curves, this effect should not cause problems with the sensor's operation. In fact, the ability to distinguish a rough finish from a smooth finish could be an advantage in some situations.

3.3.2 Measurement Repeatability

To successfully identify a material against a library of thermal response curves, the sensor's measurements should be repeatable. To test this, an aluminum block was repeatedly placed on the sensor, and the response curves were recorded. Figure 3.9 shows several of these plots. While some variation in sensor output over time




Figure 3.9: *Thermal sensor repeatability. The curves show the sensor's response to contact with the same block of aluminum during 3 trials.*

is present, in general, each of the trials produces similar results.

It should be noted that the material was placed on the sensor in a somewhat haphazard manner. For example, no special setup was used to ensure constant pressure from trial to trial. This should be similar to the sensor's actual operation conditions.

3.3.3 Variation in Heat Output

The previous experiments were performed with the sensor generating a fixed amount of heat. In particular, 0.5 Amps of current were running through the heating source, generating 0.38 watts of heat output. To determine the effects of varying heat production the sensor response at three different heat levels was monitored. A piece of aluminum was placed in contact with the sensor generating 0.09, 0.38, and 0.84 watts of heat. From the plots in Figure 3.10 one can see that increases in heat output give larger variations in the sensor's thermal response, while the general shape of the curves do not change. From this we see that increased heat production gives the sensor greater sensitivity. Of course there are limitations on the operation temperature of the sensor; for example, too high a temperature could




Figure 3.10: Thermal variation in heat output. The curves shows the temperature response profile for a block of aluminum while the heat generated by the sensor is varied. One can see that a higher heat gives the sensor a larger dynamic range.

damage the sensor or the object being sensed.

3.3.4 Thermal Temperature Recovery

When a material is removed from the sensor, some time is required for the sensor to return to its steady state temperature. To determine this recovery time, a block of aluminum was placed in contact with the sensor, and then removed. The temperature profile of the sensor was recorded during this experiment, and is shown in Figure 3.11. From this we can see that 50 seconds elapsed for a 90 percent temperature recovery. This is rather slow. However, it should be possible to reuse the sensor before full temperature recovery. In Figure 3.10 thermal profiles are shown for the same material with the sensor operating at different initial temperatures. The general shape of each of the curves is the same; the variation that occurs is just in the sensor's initial and final value. From this, we can conclude that it is not important for the sensor to operate at a constant temperature.

If faster temperature recovery was felt to be important, the heat output could be varied. By increasing the current flow through the heating layer, additional




Figure 3.11: *Thermal temperature recovery. After an object is removed from contact with the sensor a finite time is required for the pad's temperature to return to its steady state value. In the first phase of the graph, a material is in contact with the sensor. When the temperature reaches its lowest point, the object has been removed, and temperature recovery begins.*

heat could be generated on demand. In addition, the sensing scheme utilizing self heating of the thermistors would speed the recovery process. Here, the sensor attempts to maintain a constant temperature at the thermistor sites as part of the measurement process.

3.3.5 Pressure Effects

As pressure between the sensor and the sensed object is increased, the thermal contact between them also improves. This effect allows the thermal sensor to be used for tactile perception. As pressure is increased, the rate of heat flow from the sensor to the sensed object increases. Of course, this effect makes it harder to do absolute material type identification. If the contact force is unknown, or the force is changing over time, the effects due to material type will need to be isolated from the effects due to the contact pressure. This could be ascertained by giving the robot both thermal and tactile sensors.

Figure 3.12 show the sensor output for aluminum pressed into the sensor with




Figure 3.12: Thermal pressure response variation. The upper 3 curves show the response for aluminum as it is pressed into the sensor at different forces. For reference, the lower curve shows the response for brass.

3 different forces. The applied forces range from 1 to 2 pounds. For reference, the lower curve shows the response for brass. Notice that there is a significant difference between all the aluminum responses and the brass response. This indicates that even if the applied pressure is unknown, an adequate difference in sensor output exists to distinguish between those two materials.

This version of the sensor uses a covering material that attempts to minimize the effects of pressure on thermal response. The response variations detected here

are considered undesirable, since we designed the sensor for material identification. For tactile applications, a more compressible surface covering could greatly increase its pressure sensitivity.

3.3.6 Spatial Selectivity

The pixel grid of the 4×4 version of the sensor gives thermal conduction “images” of an object. This information can be used to determine a material’s spatial position, as well as its heat conduction properties. To visualize the data gathered by the array, three dimensional plots of its output are shown. The better the thermal conduction of a point, the larger its z value on the plot. In Figure 3.13 the response of the entire sensor pad as it is probed with an aluminum and wooden rod is shown.

3.4 Combining the Tactile Sensor and the Thermal Sensor

The tactile and thermal sensors previously discussed provide useful contact sensory feedback information. Ideally, a sensor could be designed that incorporates both these modalities into one device. This section briefly discusses the modifications to the tactile sensor and the thermal sensor that would be necessary to achieve this.

3.4.1 Performance Issues

Merging together the capacitive based tactile sensor and the thermistor based thermal sensor is relatively straightforward. Since the tactile sensor requires mechanical contact with the material being sensed, it should form the outer surface of the dual modality device. The thermal sensor relies upon the conductive transfer of heat, and will work when placed below the tactile device.

Stacking the tactile sensor above the thermal sensor has the undesirable effect of increasing the thermal mass of the overall device. The analysis of the thermal sensor’s performance indicates that the mass of the sensor plays a significant role in the performance and sensitivity that can be obtained. Lower sensor mass increases the rate of temperature rolloff and hence decreases the response time required for characterizing the sensed material.




Figure 3.13: Thermal spatial selectivity. Height of bar indicates level of thermal conductivity at that point. Top: contact with a round aluminum rod. Bottom: contact with both a wooden rod and an aluminum rod.

The real problem with increasing the sensor's mass is that its heat capacity also increases. To improve the performance of the tactile-thermal device, materials that minimize the overall heat capacity should be selected. To do this, the capacitive sensor should be constructed with thermally conductive components whenever possible.

9.4.2 Potential Sensor Construction

The tactile-thermal device is constructed by overlaying the thermistor based thermal sensor above the capacitive based tactile sensor. Figure 3.14 shows a cross section of the the entire device, and indicates the functionality of each surface. The lowest layer is a protective backing and mounting surface. The thermal sensor's heating supply is placed between that layer and the backing of the thermal sensor's printed circuit board. Next, the thermistors are mounted onto the top of the circuit boards, and a layer of thermally conductive silicone rubber is molded between them. The tactile sensor is mounted onto this surface. The lower conductive strips that form the force sensing capacitors are applied to this surface. Electrically conductive silicone rubber is used for the plates, and is deposited directly onto the upper surface of the thermal array. Next, the elastic-dielectric sheet is bonded onto the surface with the upper conductive traces applied to its surface. Finally, a layer of electrically insulating and then electrically conducting silicon rubber encases the sensor.

The materials used in the combined sensor should be selected according to the previous sections discussion. The elastic-dielectric material, for example, should be made of a thermally conductive material. Since this layer is rather thick compared with the rest of the device, it could form a thermal barrier, reducing the thermal sensitivity.

Since the materials to be used within this device is similar to those required for the individual tactile and thermal sensors, one would expect that its performance and ease of fabrication would also be similar. In addition, it is not much larger than either of the previous sensors, making it suitable for mounting on the Utah-MIT hand.




Figure 3.14: Tactile and thermal sensor cross section.

A Computational Architecture for Sensory Control

Chapter 4

The preceding chapters of this report have described two contact sensors suitable for use with the Utah-MIT four-fingered dexterous hand. In the near future they will be mounted onto as many as twelve surfaces of the hand's fingers, permitting detailed investigation of sensor driven control strategies. Imagine the wealth of data that will be available from these devices; clearly ample computer power to process the sensor information in real-time must be made available to fully reap the benefits that they will bring.

This chapter describes the primary components of a computational architecture with performance capabilities that meet our sensor based control requirements. Most importantly, the system provides adequate processing power, flexible software development tools, and operating system primitives appropriate for control programming.

Conventional hardware and software configurations will not be adequate for the computer requirements of the hand. Controlling a dexterous robot is a highly complex computational task in itself. The Utah-MIT hand's 16 degrees of freedom and its 32 electropneumatic actuators pose an even greater challenge; the high speed of the specially designed actuation system requires servo loop rates in

excess of 500 cycles per second. The additional computations that sensor driven control strategies will add complicates matters, and further increases the computer requirements of the overall system.

The flavor of this chapter is practical in nature; the work was initially motivated by the impending completion of the Utah-MIT hand, and the lack of a computer system powerful enough to support its control. This chapter's title, however, may be somewhat misleading in that the information is not specific to the control of the sensor equipped hand. In fact, most robotic controllers would benefit from design strategies similar to the ones described here.

The work described in this chapter is part of an ongoing development effort to provide computer hardware and software suitable for robotics real time control applications. Sundar Narasimhan has played a key role in this effort and a forthcoming joint publication will fully describe the details of the system. This chapter just presents an overview of the system's key concepts.

In the next sections, general design principles, and specific hardware and software implementation details are covered. The chapter concludes with some preliminary performance and timing results.

4.1 Design Methodology

Time was perhaps the most important constraint on the potential system's design. The intent of the project was to control a sensor equipped hand, and not to design the ultimate in computational architectures. Second to this, flexibility to accommodate various software strategies and future hardware enhancements were considered important. These two principles helped us realize a powerful system suitable for controlling the hand in a modest time period.

An early decision was made to use off-the-shelf hardware whenever possible. Custom computer hardware is usually not cost effective when only a few systems are to be built and when the computational requirements are suitable for general purpose computer architectures. In addition, keeping a custom built system at the state of the art requires constant improvements in the hardware, a never ending proposition. Since commercial companies specialize at this, it is silly for a research

group to compete with them.

An initial examination of the types of computations that are to be performed indicated that multiprocessor based hardware would work well. The computations could be partitioned onto several processors in a straightforward fashion that would not require excessive interprocessor communication.

Using multiple processors is advantageous for several reasons. Most importantly, additional computer power can be obtained incrementally by adding processors to the system. In addition, the control program structure is simplified by partitioning separate operations onto their own processor. A uniprocessor would require a more complex scheduler, since many time critical events would need to be serviced in a complex fashion. Using multiprocessors alleviates this need.

Since this system is used primarily as a research tool, it is important for the software development environment to be reasonably flexible, yet at the same time programs must run efficiently. These two goals are often conflicting, and a compromise between them must be made. An examination of two extremes in development approaches illustrates this point. The Lisp Machine provides an elegant and powerful programming environment. However its flexibility makes it unsuitable for real time performance; a program is so far removed from the machine's underlying hardware that efficiency is hard to achieve. Dedicated low performance microcomputers, such as those used in the Unimation Puma controller, are at the other end of the computational spectrum. While they are suitable for handling real time events, they must be painfully programmed in assembly language, and lack the capabilities to support any kind of reasonable development environment.

Another important question concerns the appropriateness of a computational hierarchy. In some robotics controllers, such as the MIT Puma system, only the lowest level control functions are relegated to the microprocessor front end. For many control tasks this necessitates relatively high bandwidth and low latency communications between the high level computer and the controller microprocessors. The necessary rates are often hard to achieve using conventional hardware, such as an Ethernet connection. To avoid this problem altogether, our system design assumes that almost all time critical processing, including high level functions, should be handled on the microprocessor controllers.

Table 4.1: *Computational Components of the Hand Controller*

Processor type:	Motorola 68000
Clock rate:	12 megahertz
Instruction rate:	≈ 1 MIP
Processor on board memory:	512 kilobytes
Total processors:	8
Total memory:	4 megabytes
Bus Architecture:	Intel Multibus
Analog to digital converters:	320
Digital to analog converters:	40
Host computer:	Vax via DMA Link

Finally, a key consideration in the design process was to ensure an easy path to future hardware and software upgrades. This requires all software to have a device independent structure. In fact, we are already in the process of our first system upgrade; the current Motorola 68000 processors are being converted to higher performance Motorola 68020's. So far the system port has gone smoothly, indicating that this requirement has been met.

4.2 Controller Architecture

Since almost all computer software and hardware projects seem to be given a name these days, we were forced to do so ourselves. The acronym decided upon was "MUSE". It stands for various things, as do all computer science acronyms, including the "MIT-Utah Servo loop Executor". This section describes the hardware components that the MUSE is based upon.

The processing engines used are based on a 12 megahertz Motorola 68000 single board computer. They each are equipped with 512 kilobytes of storage, a prioritize interrupt controller, three event timers, and two serial ports. These features prove useful in various software components of the controller system. For comparison purposes, they perform integer arithmetic at approximately the rate of Digital

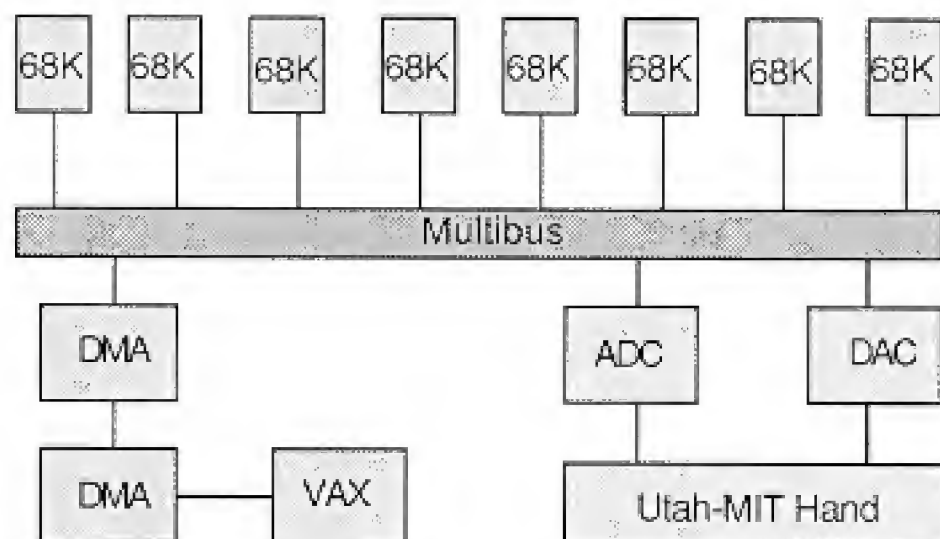


Figure 4.1: Microprocessor controller block diagram. A Multibus interconnects 8 68000 based single board computers and numerous data acquisition devices. A pair of DMA controllers link a VAX host computer to the system.

Equipment's VAX 11/750. Table 4.1 lists the system's components.

The eight processors and the peripheral equipment are interconnected with an Intel Multibus. Each processor can become a *bus master*, allowing it to take control of the bus and access the shared peripherals. A parallel oriented priority resolution controller arbitrates the processor bus requests, and grants bus control to one processor at a time in a fair fashion. In addition, each processor's memory is *dual ported*; it can be accessed locally by the processor itself, or it can be accessed globally by any other bus master. The hardware utilized provides maximum flexibility for system reconfiguration since all processors can freely access all resources on the bus. A block diagram of the system is shown in Figure 4.1.

To facilitate interprocessor communication, some additional hardware support is helpful. Since each processor's memory is dual ported communication utilizes shared memory exchanges. To insure orderly access to common state information, a test and set instruction allows semaphores to be gracefully implemented. This

instruction sets an internal processor flag and sets a bit in a specified memory location, only if the location was initially clear. The key point here is that the "testing" and "setting" occur in the same memory cycle. Without this hardware feature two processors could set a global flag at the same time, giving both of them access to the common information. The first processor could test the flag, and find it free. The other processor could obtain use of the bus on the next cycle, and also find the flag clear. The first processor would then gain control of the bus, and set the flag. The second processor would then do the same, setting the flag for the second time. Both processors would think they were the only processor with access to the shared resources.

Another feature used for interprocessor communication is a mailbox interrupt. This allows a processor to generate an interrupt on another processor by writing to a particular location in that processor's memory. The interrupted processor can then use the contents of that memory location as a flag indicating an action to take in response to the signal.

The microprocessors are connected to a VAX 11/750 host computer with a high speed *direct memory access* (DMA) link. This connection allows the VAX to write data directly into the contents of a microprocessor's memory. During a transfer the DMA controller on the Multibus becomes a master and writes directly to a microprocessor's memory through its dual port. DMA transfers from a microprocessor to the VAX are also possible. However, limitations in the structure of Multibus global interrupts make it convenient for only one microprocessor, hence referred to as the *master*, to initiate transfers in this direction.

The DMA connection between the microprocessors and the VAX allow data to be exchanged at a peak rate in excess of 200 kilobytes per second. The actual transfer rate obtained in most situations is lower and depends on such factors as the transfer size and the load on the VAX. However the bandwidth is higher than what could be achieved with either serial ports, parallel ports, or ethernet connections, since once the data transfer is initiated no processor intervention is necessary.

4.3 Software Components

Software is developed on a VAX computer running Berkeley UNIX. Code is primarily written in C, with only a few of the lowest level routines coded in assembler language. By utilizing a high level language, both the operating system primitives and the control code itself maintain a high level of hardware independence. Since the system is intended to be continuously upgraded with more advanced and powerful microprocessing hardware, software device independence is extremely important.

In a typical development scenario, programs are written on the VAX host in C, cross-compiled for the microprocessors, downloaded using the DMA link, and then executed. The compile and download cycle is rather fast due to the DMA link's speed: a download takes approximately one second. A primitive debugger that can be loaded with a program permits setting breakpoints and examining memory locations. At times the more capable debugging tools available on the VAX can also be used. Since the VAX and microprocessor have compatible C compilers, routines will often run on both systems with little modification.

A software library of useful primitive procedures allow convenient use of the hardware, and forms a key component of this system. Functions to handle inter-task communication, servo loop execution scheduling, file serving, and terminal interaction are included.

The message passing based interprocessor communication system allows the microprocessors and the host computer to exchange information. This communication scheme is device independent in the sense that the processors sending and receiving a message need not know how the information exchange actually takes place. This system is critical for building device independent application software.

The servo loop execution system provides a primitive operating system scheduler. Using the servo loop scheduler, a number of different speed servo loops can be executed on a processor at the same time. When all the servo loops are idle a background task is activated. Since servo loops are run at extremely high rates, the scheduler's efficiency is of prime concern. This component illustrates the tradeoff between efficiency and flexibility quite well. A general purpose operating system

permits maximum scheduling flexibility, at the expense of efficiency. Simple timer driven interrupt loops require little execution overhead, but they are inflexible. The servo loop scheduler is a compromise between these two conflicting requirements.

The file server allows the microprocessors to access files on the host computer. Standard Unix-like system calls are used. The file serving operations use the DMA link for high performance transfers. A process on the VAX receives file requests from the microprocessors and generates the Unix system calls to handle them. Terminal interaction among the microprocessors and the host computer utilize lower level file server capabilities. Microprocessor input and output occur over streams that are mapped to terminal input and output queues on the host.

The servo loop scheduler and the message passing system provide the basis for most of the functionality in the computational architecture support library. Together these subsystems provide the minimal level of flexibility needed for a control programming environment, and provide the control software with the necessary abstractions for device independence. The next sections discuss these systems in more detail.

4.3.1 Servo Loop Scheduling

The servo loop scheduling system (SLS) allows a processor to run various control loops at different rates in a highly efficient manner. Since a typical hand control program will have several servo loops running at rates in excess of 500 hertz, it is important for each scheduler invocation to be fast. To achieve this, scheduling flexibility has been limited to minimize the execution overhead that it requires. In fact, it is a gross overstatement to call this an operating system. It is, in fact, just an efficient utility for programming a system timer and for starting procedures based on precomputed rate information.

The program listed in Figure 4.2 uses the SLS to schedule loops running at 500 hertz and 1000 hertz. The `sls-schedule` routine is called for each task; a symbolic name for the procedure, the C procedure name itself, and its rate in hertz, are passed as arguments. The call returns a pointer to the servo loops scheduling parameters. The `sls-start` initiates execution of the loops. The code


```
main()
{
    int force, position;

    /* initialize the system */
    force = sls-schedule("force control loop", servo1(), 1000);
    position = sls-schedule("position control loop", servo2(), 500);

    /* start in position control mode */
    sls-pause(force);
    sls-start();

    /* the following code runs as the "background" job */
    sls-pause(position); /* now switch to...
    sls-resume(force)    ; /* ... force control mode */

    sls-pause(force);    /* now switch to...          */
    sls-resume(position); /* ... position control mode */
}

servo1()
{
    /* code for the force control loop */
    force_control_update();
}

servo2()
{
    /* code for the position control loop */
    position_control_update();
}
```

Figure 4.2: *Sample use of the servo loop scheduler. In this example, a program schedules two servo loops. In the background job code, force control and position control are selected using the `sls-pause` and `sls-resume` commands*

following the `sls-start` command is referred to as the *background job*, and runs when the servo loops are idle.

Two additional commands, `sls-pause` and `sls-resume`, allow specified servo loops to be suspended and restarted. These commands can be used by a servo loop or the background job. As shown in the example, a system that has a position control mode and a force control mode might schedule two loops; one for position control, and the other for force control. The `sls-pause` command could be used to disable the position control loop when the system is running in force control mode. To switch to position control, the force control loop would be paused, and the position control loop resumed.

A benefit of the SLS is that it simplifies the structure of control programs; the details of the interrupt and timer hardware are not mixed in with the control program code. This increases the overall portability of the system.

To minimize execution overhead, the SLS is table driven. An *event table* is automatically generated by the system when the `sls-start` command is issued. This table lists the elapsed time between invocations of the scheduled servo loops. For example, the event table for two servo loops, one running every ten seconds and the other running every five seconds has two entries. The first entry indicates that both loops are to start, and five seconds elapse until the next event. The second entry indicates that the five second loop should start, and another five seconds are to elapse before the next event. After this, the cycle repeats, and the first entry of the event table is reused.

With the system outlined so far, it is possible for more than one loop to be runnable at the same time. The system must have an orderly method for selecting the actual loop that will be run from the set of runnable loops. A *process table* is maintained for this purpose. All the tasks in the system are arranged, in order of decreasing servo rate, in this table. When the event table indicates a loop is ready to run, it is marked runnable in the process table. The system then searches down the process table, and starts the fastest rate loop that is marked runnable.

The time to the next event stored in the event table is loaded into a timer on the processor. When the time has elapsed, the running task is interrupted, and the scheduler is reinvoked. The next tasks in the event table that are scheduled to

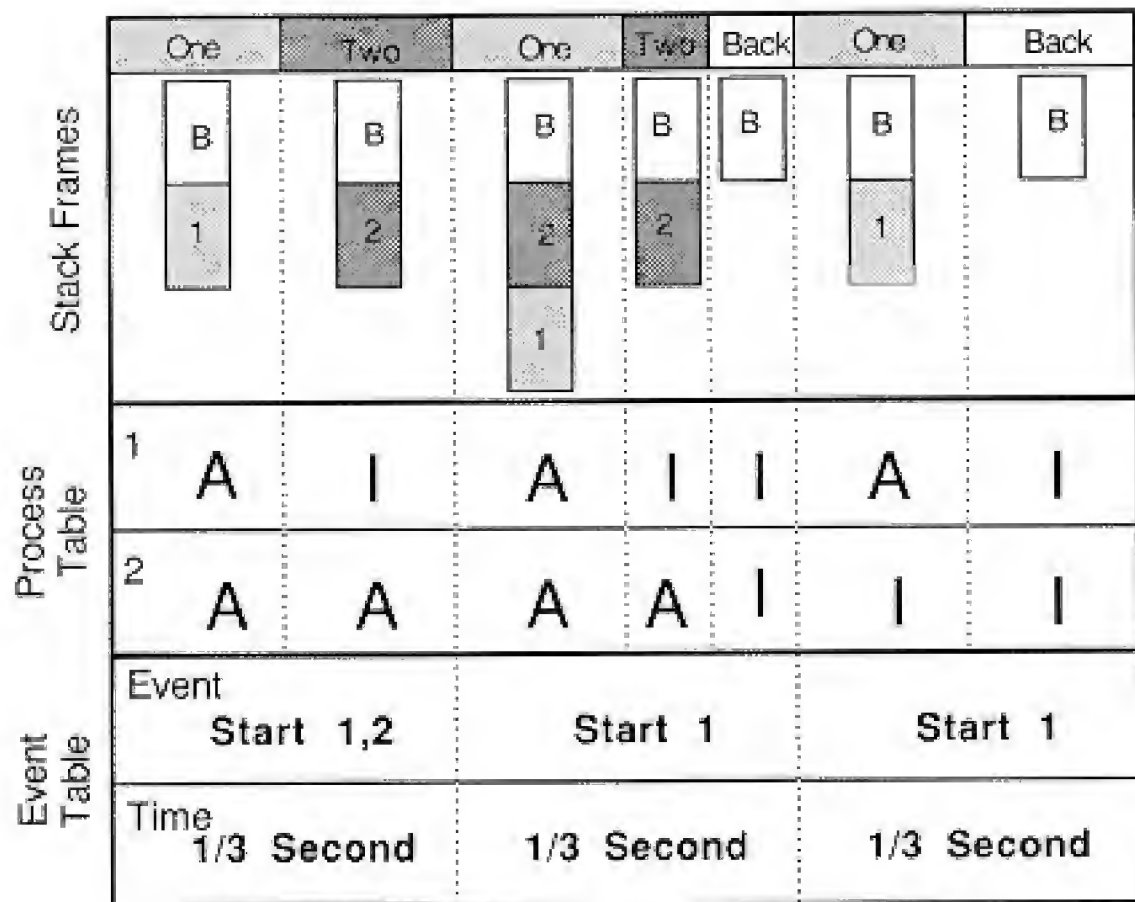
start are marked runnable in the process table. If a loop with a speed slower than the interrupted loop is made runnable, the interrupted loop will be resumed. If a higher speed servo loop is made runnable the slower loop that was interrupted will be temporarily suspended, until higher speed loops that are runnable complete.

An implication of assigning a priority to a process based on its rate is that a loop can only be interrupted by a higher speed loop, and hence, no coprocessesing can take place. This is not considered to be a problem. The rate specified for a servo loop is a request that the loop be run that number of times a second. The exact time that a loop is invoked is not important, as long as it is runs within its specified time slice. In other words, a loop scheduled to run every second is only a guarantee that the loop will run *sometime* within a second. A finer precision in selecting the time at which a procedure will run is not needed within our control programming scenario.

Coprocessing is not supported in the SLS's restricted scheduling environment. Eliminating coprocessing results in a convenient simplification to the system; only one stack need be maintained for all the servo processes running on a processor. Stack pointers are not changed when a new process is invoked, or a suspended process is resumed.

When a loop terminates, the scheduler is also invoked. The terminating loop is marked idle in the process table, and a new loop is selected to run. If no servo loops in the process table are runnable, the background job is activated.

To help clarify the above discussion, Figure 4.3 diagrams the relationship between the different components of the SLS. In this example, two loops are scheduled. Loop one is invoked three times a second, and loop two is invoked once a second. The first event table entry sets loops one and two to runnable in the process table. The system then selects loop one to run, since it runs at a faster rate than loop two. Loop two will not run until it is the fastest loop marked active in the process table. The next event occurs 1/3 of a second later, when loop one is to be reinvoked. We see that in the first time slice of the diagram, loop one is active, and the system stack frame has the activation record for loop one on its top. When loop one terminates, the scheduler starts the fastest active loop in the process table, in this



A: Active
I: Idle

Figure 4.3: Servo loop scheduler timing diagram. The relationship between running servo loops, elapsed time, the system stack, the process table and the event table are diagrammed. Two loops are scheduled, one running three times a second and the other running once a second.

case, loop two. While loop two is running, the next event occurs; the reactivation of loop one. Since loop one is now the fastest loop marked active in the process table, it will interrupt loop two. This occurs in time slice three of the diagram. The stack frame now has the activation record for loop one on top of the record for loop two. When loop one terminates, its record is popped off the stack, and loop two is resumed. This occurs in time slice four. In time slice five loop two

terminates, none of the servo loops are active, and the background job is resumed. Finally, the third event in the cycle occurs, and loop one is reinvoked. After this, the process repeats with the first entry in the event table.

It is easy to detect when a processor cannot maintain a given set of servo loop rates. If a clock tick occurs and a servo loop that the event table indicates is ready to run, is already marked runnable in the process table, the desired servo rate cannot be maintained. In some cases we allow servo loops to overrun a small percentage of the time. For example, an extraordinary event that may require extra processing in a particular invocation of a loop can cause an infrequent, but tolerable, overrun.

In timing tests it was found that the raw overhead of invoking a procedure with the servo loop scheduler is low. In an application with one servo loop, 39 microseconds were required to invoke the scheduler, on the average. In a test with four loops being scheduled to run at the same rate, 24 microseconds were required to schedule each loop. In a case with eight loops being scheduled at the same rate, 17 microseconds were needed. The decrease in average time spent within the scheduler is due to the low cost of restarting a queued process. When many loops are scheduled at the same rate, they all are *readied* on the same clock tick. When a loop finishes, the time required to start a readied, but suspended, loop is very low. Another way of looking at this is that scheduling a set of processes at rates that minimize interrupt overhead is most efficient.

4.3.2 Intertask Communications

In a typical control application there are a number of tasks being scheduled by the SLS that must communicate with one another. The tasks may be running on any of the micros or host computers. Without a transparent form of communication, each task would need the exact routing information needed for communicating with another routine, forcing the programmer to wire the execution structure of the overall system into the control programs themselves. To make matters worse, the task would need to use different methods of communication depending on the specific processor receiving the information. To overcome these difficulties,

Table 4.2: The Fields of a Message Buffer

Field	Explanation
<code>sender</code>	Virtual processor ID that sent message.
<code>recipient</code>	Virtual processor ID to receive message.
<code>sender-vdd</code>	Virtual mailbox for a reply message.
<code>recipient-vdd</code>	Virtual mailbox to receive message.
<code>message-id</code>	One byte user assigned message type.
<code>message-free</code>	Flag indicating message buffer is free.
<code>buffer-length</code>	Message buffer length.
<code>buffer</code>	Message buffer.

a message passing communication system is utilized. A uniform communication protocol sends data between tasks running on the same processor, and between tasks running on different microprocessors or a host computer.

Tasks send and receive data through uniformly formatted messages. The fields in a message buffer can be classified into two broad categories. The *fixed format* fields are required, and are used by the message passing system for specific purposes. The *free format* message buffer is optional, and is used to implement higher level protocols using the low level transport mechanism.

Messages are allocated with the `balloc` routine and are deallocated using `bfree`. In general, the messages are not freed by the processor that allocated them, and hence `bfree` cannot actually release the storage. Instead, `bfree` just sets a garbage collection bit in the message buffer. At a later time the sending processor that issued the `balloc` garbage collects the freed buffers. In the current implementation calls to `balloc` periodically invoke the garbage collector, although other schemes could be employed.

A message contains information that allows the router to send it to a recipient process, and information that allows the recipient process to send back a reply. To be more specific, a recipient has two parts: the sender virtual processor, and a routine that runs on that virtual processor. Most commonly, virtual proces-

Table 4.3: *Routing table example. This system has 4 virtual processors running on 3 real processors.*

Processor Zero			Processor One		
Processor	Channel	Address	Processor	Channel	Address
zero-a	self	jump-table-a	zero-a	mbus	zero
zero-b	self	jump-table-b	zero-b	mbus	zero
one	mbus	one	one	self	jump-table-a
VAX	DMA		VAX	mbus	zero

Processor VAX		
Processor	Channel	Address
zero-a	DMA	
zero-b	DMA	
one	DMA	
VAX	self	jump-table-a

sors correspond to the real processors in the system, and the recipient is just the processor number and a routine on that processor to invoke. The message buffer fields `sender` and `recipient` contain the virtual processor number of the sending processor and recipient processor respectively (see Table 4.2).

The recipient routine that is invoked when a message arrives is referred to as a *virtual device driver* (VDD). The message field `recipient-vdd` contains an offset in a table on the receiving processor that points to the address of the routine that is invoked for that message. Reply messages that are directed back to the sender invoke the routine indicated by `sender-vdd`.

The table associating VDD numbers to routine addresses is referred to as the *jump table*. Every virtual processor has its own jump table. The `sender-vdd` and `receipt-vdd` fields contain offsets into this table, and not actual routine addresses, to simplify building a message passing system that runs on multiple processors. Since addresses are not assigned to a procedure until a program is linked, it is easier to refer to the VDD symbolically. The additional overhead that this requires

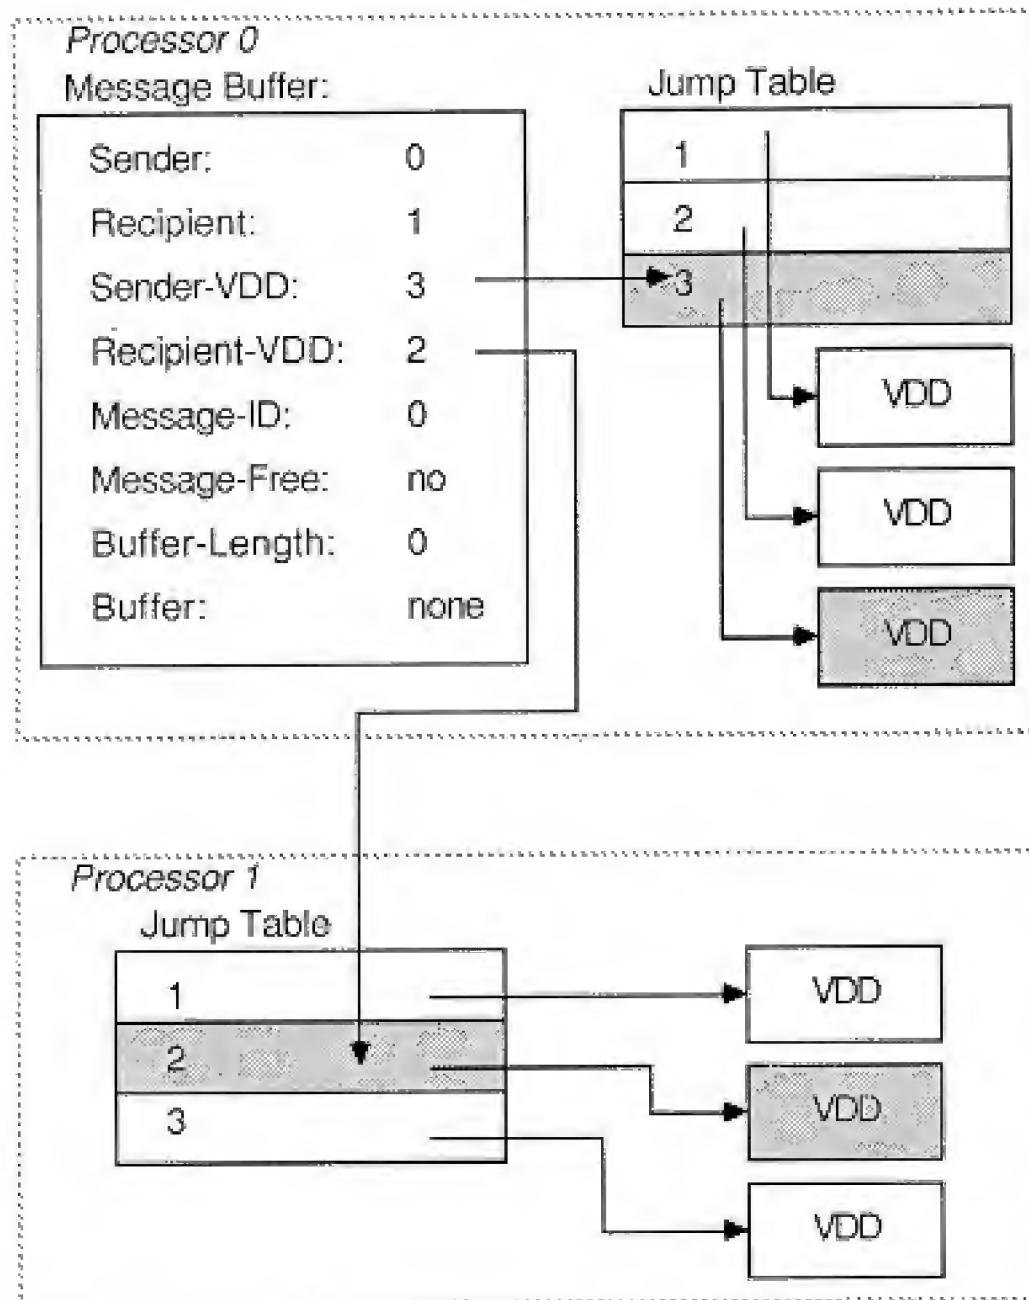


Figure 4.4: Message passing system components. The VDD fields point to an offset in a jump table on the recipient processor.

is small, and considered acceptable.

Figure 4.4 diagrams the relationships among the different components of the message passing system. Here, the message buffer field `sender-vdd` contains an offset in the recipient processors jump table. The jump table contains the addresses of the routines that correspond to the VDDs assigned to that processor. Notice that the sending processor has no information as to the address of the VDD routine, allowing dynamic modification of VDD handler assignments. The `sender-vdd` field points to an offset in the sender's jump table. This VDD offset is used by the recipient processor to reply to the message.

The actual process of sending a message is controlled by one routing table on each physical processor. This table associates a virtual processor name with a communication channel and an address on that channel. A communication channel corresponds to a routine that is invoked to send the message. The address field is passed as an argument to that routine. In general, each physical method of sending a message has one communication channel routine. For example, the Multibus channel sends messages between microprocessors connected over that bus. The DMA channel sends messages between the VAX host processor and microprocessor zero. A pseudo channel called “self” sends local messages, and in effect causes a branch to the associated jump table for execution of the message handler.

Table 4.3 shows the routing tables for a message passing system with four virtual processors (zero-a, zero-b, one, and VAX) that run on three real processors (zero, one, and VAX). In this system, only real processor zero can communicate with the VAX using the DMA channel. Message forwarding must be used for processor one and processor VAX to communicate. In the routing table for processor one, messages destined to the VAX are sent to processor zero. When processor zero receives this message, the router notices it is intended for processor VAX, and uses the DMA channel to complete the transfer. An arbitrary number of message forwarding hops can be handled by this scheme.

A typical message passing system is shown in Figure 4.5. Two VDDs and three processors have been defined. The `serveVDD` mailbox runs on processors zero and one while the `vaxdataVDD` runs on processor VAX. The main routine, possibly

```

main()
{
    message = balloc();
    message->buffer = force_information;
    sendmessage(servoVDD,message,0);
    sendmessage(servoVDD,message,1);
}

servoVDD(message)
{
    /* process the message */
    process_servo_data(message->buffer);

    /* send data to the VAX */
    another_message = balloc();
    another_message->buffer = random_data;
    sendmessage(vaxdataVDD,another_message,V);

    /* free the buffers when done */
    bfree(message); bfree(another_message);
}

vaxdataVDD(message)
{
    process_vax_data(message->buffer);
    bfree(message);
}

/* system wide message passing specification file */
$VDD(servoVDD, 01)
$VDD(vaxdataVDD, V)
$PROCESSOR(PROC0,0)
$PROCESSOR(PROC1,1)
$PROCESSOR(PROCVAX,VAX)

```

Figure 4.5: Typical message passing system specification. Two VDDs and three processors have been defined.

a master control program, sends a message to the `servoVDD` mailboxes on both processor zero and processor one. Each `servoVDD` receives the message and then forwards some data to the `vaxdataVDD` that runs on the VAX. The message passing system specification file, shown at the bottom of the figure, is used by the system to build the routing and jump table structures used by this configuration.

A major advantage of using virtual device drivers is the ease in which they can be relocated. For example, a VDD to invert matrices might be located on the VAX. If it is later decided to use a microprocessor to handle the matrix inversions, we just update the `@VDD` commands to indicate their new location and recompile the system.

The cost of the flexibility provided is minimal. All VDD routing is fixed at compile time, and does not incur substantial runtime overhead. In fact, the entire system is table driven. When a message arrives, the message handler issues a jump to the proper VDD handler. No procedure call overhead is required to process the message. Messages exchanged between the VAX and the micros are also relatively fast, due to the speed of the DMA link.

The speed of the message passing system is demonstrated by a set of benchmark tests performed. Sending a message from one microprocessor to another takes 20 microseconds. Sending a message and receiving a return reply takes 40 microseconds. In comparison, Digital's RSX operating system's message passing is slower; 20 milliseconds are required to send a message, and 60 milliseconds elapse when sending a message and receiving a reply. The time to send a message from a microprocessor to the VAX, in our system, is slower. It takes 1 millisecond to send to the hand processor. This still is faster than the typical time required to send an RSX message. It should be noted that the speed advantage our system has over RSX is gained by simplicity. We feel, however, that in the domain of control programming no functionality has been lost.

4.4 Performance Evaluation

The computational architecture described is now operational and is actively being used to control the Utah-MIT hand. Initially 5 microprocessors were use in the

system. One processor was dedicated to each finger, and the remaining one was the master coordinating inter-finger activities and communications with the VAX. However, the computational power provided proved inadequate for some of the initial tasks we wanted to perform. Each 68000 provides approximately 1 MIPS performance. If each processor serves 4 joints at 500 hertz, approximately 500 integer instructions per joint per servo update are available. The C code used rather quickly exceeded this instruction quota.

The system was then updated to 8 68000 processors, increasing overall performance by 3 MIPS. The computations were partitioned somewhat differently: one processor was assigned all data acquisition needs, 5 processors handled finger computations, and 1 processor handled inter-finger coordination. This leaves one spare computer for future expansion. This level of performance has been adequate for the current tasks being performed; full digital positional control and force control have been obtained. Floating point operations are not used in the current control software. Instead, integer arithmetic with appropriate scaling operations is employed.

The current computations being performed are far less complex than those planned for the near future. Grasping operations will require computing the equivalent of a grip jacobian [Mason and Salisbury 1985; Hollerbach et al. 1986], a relatively complex task. In addition, the current lack of hardware floating point support may prove to be a problem. These computations will probably require the use of floating point operations to avoid numerical instabilities. The current hardware cannot handle these needs. Their lack of floating point support is particularly troubling; performing a floating point operation in software can often take 50 times longer than it would take using specialized hardware.

The issue of whether floating point support is really required for control computations is often debated. Using integer arithmetic with proper scaling can often substitute for a full floating point implementation. However, certain computations cannot always be scaled by a fixed amount, requiring more complex dynamic scaling procedures. In the end, since the effort required to isolate the proper scaling operations may be large, and since dynamic scaling becomes quite complex, having hardware floating point capabilities is probably warranted.

To meet these needs, the computational architecture is currently being upgraded to Motorola 68020 processors based on the VME bus. The 68020 itself gives a factor of 2.5 speed improvement over the 68000 in integer arithmetic. In addition, the processor supports the Motorola 68881 floating point coprocessor. The upgraded system will employ 10 processors giving a total of 25 MIPS integer performance, a significant increase over the 8 MIPS available on the current system. Sun workstations, also based on the Motorola 68020 architecture, will be used for development. Control programming will be simplified by using the same processor for both the microprocessors and the host computer.

To port the system, only limited modification to a few parts of the software library are required. The device dependent portions of the code are limited to low level software such as interrupt handlers, and timer and serial port controllers. All higher level control code is entirely device independent and need just be recompiled.

Conclusions

The opening paragraphs of this thesis reflected on the amazing capabilities of our hands. No robots have yet been able to achieve their level of versatility in manipulatory and exploratory motions. Our hands can successfully function in unstructured environments that foil even the most sophisticated manipulators built to date. The most important shortcomings occur with their dexterity, sensing capabilities, computational power, and algorithmic support.

This thesis has addressed current robots' deficiencies in sensors and computational support. The preceeding chapters discussed a tactile sensor, a thermal sensor, and a computational architecture suitable for use with an advanced dexterous robotic device. The tools needed for realizing a robot suitable for use in an unstructured environment are now in place; the Utah-MIT hand provides a high level of dexterity, the sensors provide environment feedback, and the computational architecture gives it the needed number crunching power.

The next sections in this concluding chapter briefly summarize the work presented in this thesis, and then discuss future research directions to be pursued.

5.1 Where Have We Gone So Far?

A tactile sensor and a thermal sensor have been design and tested. The technologies employed permit fabrication of highly durable devices that can be mounted on curved surfaces like the fingers of the Utah-MIT hand. Their performance was found to be adequate for some of the sensor driven control tasks that are planned to be implemented on the Utah-MIT hand.

The tactile array provides 8×8 forces sensors with 1.9 mm center-to-center spacing. Each cell provides 8 bits of force data over a 0 to 200 gram per square mm range. Its actual sensitivity can be controlled by varying the properties of its elastic-dielectric layer. The device's overall performance is influenced by the size of the rubber tabs, the thickness of the covering sheet, and the actual silicone rubber used. Although the capacitance being detected is small, proper electrical shielding minimizes the noise created by parasitic affects. Hysteresis was found to be low, though a discussion questioned the importance of achieving low hysteresis; since the fingers need soft coverings, it is unlikely that hysteresis can altogether be eliminated. Importantly, human touch sensing has a high level of hysteresis yet it maintains excellent performance.

The thermal device has a matrix of 4×4 cells with 3.5 mm center-to-center spacing. By heating an object and measuring the resulting temperature change at the sensor's surface, a material's type can be determined. Since the sensor has an array of thermal cells, it can also detect the contact outline of an object. Several different materials were classified using their characteristic thermal profile.

The feasibility of combining the tactile and thermal sensor into one package was covered. Since the thermal sensor is passive, and does not require direct mechanical contact with the external object being sensed, a layered device can be constructed with the thermal sensor placed below the tactile sensor. The choice of materials is crucial for the performance of this dual modality device; the tactile sensor must be made of thermally conductive materials to allow heat to be exchanged freely with the thermal sensor.

Since a sensor equipped dexterous robot will require a high degree of intelligence, and since the computational tasks that must be performed are likely to be

complex, a computational architecture with suitable performance was developed that meet these needs. The system described is actively being used to control the Utah-MIT hand. It performs low level position and force control servos at rates high enough for adequate performance [Hollerbach and Narasimhan 1986]. The hand control strategies being used have ample parallelism to be partitioned onto a set of 8 processors. The architecture is general purpose enough to permit use with other robotics control projects. Currently, the MIT direct drive arm utilizes a system based on the components of this computational architecture.

5.2 What Comes Next?

This section explores some future areas of research in sensing and computational support that will be required to fully realize a dexterous hand that can operate in unstructured environments. The discussion is divided into two sections. The first covers the remaining work on the tactile and thermal sensors necessary for operating them with the Utah-MIT hand. The second discusses some higher level issues involved with achieving sensor based control.

5.2.1 *Improving the Tactile and Thermal Sensors*

Some work remains to be done before a fully operational sensor equipped dexterous hand can be realized. The final details of the tactile sensor's fabrication process must be worked out to insure long term reliability. A reduction in the size of a tactile cell to 1.2 mm is under strong consideration. The current 1.9 mm center-to-center spacing is somewhat large for some planned tasks. Fortunately, the junction capacitance available for the smaller sized cells will still be large enough to permit detection with adequate force resolution.

Bonding wires to the flexible pad is another important consideration. Up to now the sensor has been mounted on a rigid printed circuit board. In the near future, fully flexible versions will be made that permit mounting on the curved surfaces of the fingers. Small wires must come out from the pad and be attached to the detection and scanning electronics. It is important for these connections to be extremely reliable. An alternative design that uses a flexible printed circuit

board that can be wrapped around the hand's fingers is also under development. This approach might lead to more reliable attachments.

Since current plans call for mounting 3 tactile arrays on each of the hand's 4 fingers, some work must be done to further miniaturize the scanning and detection electronics. Currently the only electronics mounted near the sensor array are the primary row amplifiers. They must be placed near the pad to reduce stray capacitance effects. To minimize the wiring nightmare that the 12 sensor pads would create, multiplexing must be done on the fingers themselves. Since the space available is so small, hybrid circuit fabrication is being considered. This will also permit further signal amplification to be done close to the sensors, possibly increasing their overall performance.

More attention must be given to modeling the tactile sensor's performance. A detailed mechanical model would give a better understanding of many of the design decisions that were made. For example, the elastic-dielectric material has many design parameters that can be varied. The actual mechanical properties of the silicone rubber, the size and height of the rubber tabs, and the thickness of the surface covering can all be varied. A detailed analysis of how these factors affect overall sensor performance is certainly warranted. Alternatives to the rubber tabs could be considered. A continuous silicone rubber material with embedded air pockets, like a sponge, is also a possibility.¹

A sensor's reliability is of the utmost importance. Unfortunately, putting a sensor in a durable package often reduces its sensitivity. The sensors discussed in this thesis are fairly reliable, though much more testing is needed to ascertain their long term performance. Perhaps it would be best to form the sensor with an easily replaceable covering. Provisions must be made, however, to allow the covering to bond firmly to the sensor's surface, or bad mechanical interactions between them might result.

Some work remains to be done before the combination tactile and thermal sensor can be realized. Since the thermal sensor is to be placed behind the tactile sensor,

¹An advantage of the tabs is that they concentrate the rubber at the capacitor's junction and increase its dielectric constant. A sponge like material would have some air pockets at junction sites, reducing the dielectric effect.

the heat capacity and heat conductivity of the elastic-dielectric material must be suitably high. If the tactile sensor formed a thermal barrier, the performance of the thermal sensor would be seriously degraded. Materials suitable for this purpose must be isolated.

5.2.2 Sensor Based Hand Motions

The broader issues of how to integrate contact sensor data into hand control strategies needs further investigation. To date, few successful system that actively utilize contact sensor data have been developed. With robots like the Utah-MIT hand and sensors like the ones described in this thesis, advances in this field should now be possible.

A fundamental analysis of tactile sensor based operations must be performed. Take the task of programming the hand to grip objects and to automatically compensate for potential slippage. Firstly, the phenomena of slippage must be thoroughly investigated. The exact mechanical interactions between the hand and the object, as seen by the fingertip contact sensors, must be quantified. Then the grasp planner must be designed to utilize not only the geometrical model of the object and the hand, but the sensing data being acquired in real time. In essence, the constraints provided by sensory data must be considered with the same importance as those provided by the static world model.

Much information can be gained by investigating how the human hand uses sensory information. Lederman and Klatzky's [1985] pioneering work in the area of haptic perception and recognition of common objects is a useful start. They define a set of "knowledge directed procedures" that the haptic system might use to obtain information about a material's intrinsic properties. For example, to obtain an object's weight, an unsupported holding strategy is used. Shape and size can be found by enclosure or dynamic contour exploration. Textural information is obtained by relative motions between the hand and the object.

The Lederman and Klatzky work emphasizes the importance of the hand as a sensory organ. A robotic hand can be made to perform some of these same haptic functions. The sensory data it obtains can supplement other sources of

information, primarily vision. The whole area of haptics and fusion of multiple sensory sources requires much additional research. Ultimately, a robust robotics system that can handle large amounts of modeling uncertainty will be possible. In fact, such a robotics device will actually be able to explore its own environment, obtaining the information that it needs to accomplish particular operations.

Are there tasks that a hand with touch sensors can perform that a hand without sensors is incapable of achieving? This question is asked by many, though no formal answer has yet been given. Most researchers in the field believe that contact sensory information plays a crucial role in the human hand's operations, and conclude that the same holds for robot hands. The robotic sensor field would be given a big boost by a demonstration that this assumption is really correct. Perhaps the most convincing argument would be a demonstration of a sensor equipped robotic hand performing a complex operation that to date has only been done by humans.

Though much work remains, this thesis has shown that an artificial hand with advanced sensory capabilities is now well within our reach. The sensor equipped Utah-MIT hand has the potential to provide advances in automation and prosthetic devices, and to give us a better understanding of motor control. Devices like this will allow robots to emerge from the structured world that they are confined to today.

References

1. Abramowitz, J., Goodnow, J., and Paul, B., "The Pennsylvania Articulated Mechanical Hand," *Proc. ASME Conference on Robotics*, Chicago, 1983.
2. Baird, H. S., Wells, E. G., Britton, D. E., "Coordination Software for Robotic Workcells," *Proc. IEEE International Conference on Robotics and Automation*, Atlanta, pp. 354-360, 1984.
3. Bajcsy, R., Brown, D., Wolfeld, J., Peters, D., "What Can We Learn From One Finger Experiments?," *Computer and Information Science Department*, University of Pennsylvania, PA, 1982.
4. Begej, S., "An Optical Tactile Array Sensor," *Proc. SPIE Conference on Intelligent Robots and Computer Vision*, pp. 271-280, Cambridge, MA, 1984.
5. Biggers, G. E., Jacobsen, S. C., K. B., Gerpheide, "Low-level Control of the Utah-MIT Dexterous Hand," *Proc. IEEE International Conference on Robotics and Automation*, San Fransisco, pp. 61-66, 1986.
6. Binford, T. O., "Sensor Systems for Manipulators," *Proc. of Conference*

- on Remotely Manned Systems*, Cal Tech, pp. 283-291, 1972.
7. Boie, R. A., "Capacitive Impedance Readout Tactile Image Sensor," *Proc. IEEE International Conference on Robotics and Automation*, pp. 370-379, Atlanta, 1984.
 8. Brockett, R. W., "Robot Hands with Rheological Surfaces," *Proc. IEEE International Conference on Robotics and Automation*, pp. 942-947, St. Louis, 1985.
 9. Bruno, G., Demartini, C., Valenzano, A., "Communication and Programming Issues in Robotic Manufacturing Cells," *Proc. IEEE International Conference on Robotics and Automation*, Atlanta, pp. 361-367, 1984.
 10. Butcher, A., Fehrenbach, P., "A Computer-controlled reconfigurable gripper," *General Electric Company, Technical Report*, 1983.
 11. Carslaw, H. S., Jaeger, J. C., "Conduction of Heat in Solids," Oxford at the Clarendon Press, Great Britain, 1948.
 12. Chomerics Corporation, "EMI Shielding Engineering Handbook," Woburn, MA, 1985.
 13. Chun, K. J., Wise, K. D., "A Capacitive Silicon Tactile Imaging Array," *Proc. IEEE International Conference on Solid-State Sensors and Actuators*, pp. 22-25, Philadelphia, 1985.
 14. Dario, P., De Rossi, D., Domenici, C., Francesconi, R., "Ferroelectric Polymer Tactile Sensors with Anthropomorphic Features," *Proc. IEEE International Conference on Robotics and Automation*, pp. 332-340, Atlanta, 1984.
 15. Dario, P., Bicchi, A., Vivaldi, F., Pinotti, P. C., "Tendon Actuated Exploratory Finger with Polymeric Skin-Like Tactile Sensor," *Proc. IEEE International Conference on Robotics and Automation*, pp. 701-706, St. Louis, 1985.

16. Dario, P., De Rossi, D., "Tactile Sensors and the Gripping Challenge," *IEEE Spectrum*, vol. 22, no. 8, August 1985.
17. Dupourque, V., "A Robot Operating System," *Proc. IEEE International Conference on Robotics and Automation*, Atlanta, pp. 324-348, 1984.
18. Fearing, R. S., Hollerbach, J. M., "Basic Solid Mechanics for Tactile Sensing," *Proc. IEEE International Conference on Robotics and Automation*, pp. 266-275, Atlanta, 1984.
19. Goldwasser, S. M., "Computer Architecture for Grasping," *Proc. IEEE International Conference on Robotics and Automation*, Atlanta, pp. 320-325, 1984.
20. Grahn, A. R., Astle, L., "Robotic Ultrasonic Force Sensor Arrays," *Robotic Sensors*, vol. 2, pp. 297-315, IFS Publications, New York, 1986.
21. Grimson, W. E. L., Lozano-Perez, T., "Model-Based Recognition and Localization from Tactile Data," *Proc. IEEE International Conference on Robotics and Automation*, Atlanta, pp. 248-255, 1984.
22. Hackwood, S., Beni, G., Hornak, L. A., Wolfe, R., Nelson, T. J., "Torque-Sensitive Tactile Array for Robotics," *The International Journal of Robotics Research*, vol. 2, no. 2, pp. 46-50, 1983.
23. Harmon, L. D., "Touch-Sensing Technology: a Review," Tech. Rep. MSR 80-83, Society of Manufacturing Engineers, Dearborn, Michigan, 1980.
24. Harmon, L. D., "Automated Tactile Sensing," *The International Journal of Robotics Research*, vol. 1, no. 2, pp. 3-32, 1982.
25. Hillis, W. D., "A High-Resolution Imaging Touch Sensor," *The International Journal of Robotics Research*, vol. 1, no. 2, pp. 33-44, 1982.
26. Hollerbach, J. M., "Workshop on the Design and Control of Dexterous Hands," *AI Lab Memo no. 661*, Massachusetts Institute of Technology, April 1982.

27. Hollerbach, J. M., "Tactile Sensors and Interpretation of Contact Features," *Proceedings of the NSF Workshop on Intelligent Robots: Issues and Achievements*, SRI International, Menlo Park, CA, Nov. 13-14, 1984.
28. Hollerbach, J. M., Narasimhan, S., Wood, J. E., "Finger Force Computation without the Grip Jacobian," *Proc. IEEE International Conference on Robotics and Automation*, San Francisco, pp. 871-875, 1986.
29. Horowitz, P., Hill, W., "The Art of Electronics," Cambridge University Press, Cambridge, 1980.
30. Ikegami, A., Kaneyasu, M., "Olfactory Detection Using Integrated Sensor," *Proc. IEEE International Conference on Solid-State Sensors and Actuators*, pp. 136-139, Philadelphia, 1985.
31. Jacobsen, S. C., Wood, J. E., Knutti, D. F., Biggers, K. B., "The Utah-MIT Dextrous Hand: Work in Progress," *The International Journal of Robotics Research*, vol. 3, no. 4, pp. 21-50, 1984.
32. Klein, C. A., Wahawisan, W., "Use of a Multiprocessor for Control of a Robotic System," *The International Journal of Robotics Research*, vol. 1, no. 2, pp. 45-59, 1982.
33. Lecombe, M. H. E., "Tactile Sensors, Sonar Sensors, and Parallax Sensors for Robot Applications," *Proc. 6th International Symposium on Industrial Robots*, Univ. Nottingham, U.K., pp. 22-26, 1976.
34. Lederman, S., Klatzky, R., "Hand Movements: A Window into Haptic Object Recognition," *Proc. of the 26th Annual Meeting of the Psychonomic Society*, Boston, 1985.
35. Mason, M. T., Salisbury, J. K., "Robot Hands and the Mechanics of Manipulation," MIT Press, Cambridge, 1985.
36. Narasimhan, S., Siegel, D. M., Biggers, K., Gerpheide, G., Hollerbach, J. M., "Implementation of Control Methodologies on the Compu-

-
- tational Architecture for the Utah-MIT Hand," *Proc. IEEE International Conference on Robotics and Automation*, San Francisco, pp. 1884-1889, 1986.
37. Nomura, A., Ichimatsu, A., Shibata, I., Watanabe, I., Nihei, K., "Two Dimensional Tactile Sensor Using Optical Methods," *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, pp. 264-268, CHMT-8, 1985.
38. Ogorek, M., "Tactile Sensors," *Manufacturing Engineering*, vol. 94 no. 2, pp. 69-77, 1985.
39. Okada, T., "Object Handling system for Manual Industry," *IEEE Transactions on System, Man and Cybernetics*, vol. SMC 9, no. 2, pp. 79-89, February 1979.
40. Okada, T., "Computer control of Multijointed finger system for precise object handling," *IEEE Transactions on System, Man and Cybernetics*, vol. SMC 12, no. 3, pp. 289-299, May-June 1982.
41. Overton, K. J., Williams, T., "Tactile Sensation for Robots," *Proc. 7th International Joint Conference on Artificial Intelligence*, pp. 791-795, Vancouver, Canada, 1981.
42. Petersen, K., Kowalski, C., Brown, J., "A force Sensing Chip Designed for Robotic and Manufacturing Automation Applications," *Proc. IEEE International Conference on Solid-State Sensors and Actuators*, pp. 30-32, Philadelphia, 1985.
43. Phillips, J. R., Johnson, K. O., "Tactile Spatial Resolution. III. A Continuum Mechanics Model of Skin Predicting Mechanoreceptor Responses to Bars, Edges, and Gratings," *Journal Of Neurophysiology*, vol. 46, no. 6, pp. 1204-1225, Dec. 1981.
44. Purbrick, J. A., "A Force Transducer Employing Conductive Silicone Rubber," *Proc. of the 1st International Conference on Robot Vision and Sensory Controls*, pp. 73-80, Stratford-upon-Avon, UK, 1981.

45. Raibert, M. H., Tanner, J. E., "Design and Implementation of a VLSI Tactile Sensing Computer," *The International Journal of Robotics Research*, vol. 1, no. 3, pp. 3-18, 1982.
46. Raibert, M. H., "An All Digital VLSI Tactile Array Sensor," *Proc. IEEE International Conference on Robotics and Automation*, pp. 314-319 Atlanta, 1984.
47. Russell, R. A., "A Thermal Sensor Array to Provide Tactile Feedback for Robots," *International Journal of Robotics Research*, vol. 5, no. 3, pp. 35-39, 1985.
48. Rohsenow, W. M., Choi, H., "Heat, Mass and Momentum Transfer," Prentice-Hall, Englewood Cliffs, NJ, 1961.
49. Salisbury, J. K., "Design and Control of an Articulated Hand," *Proc. International Symposium on design and synthesis*, Tokyo, 1984.
50. Schneiter, J. L., Sheridan, T. B., "An Optical Tactile Sensor for Manipulators," *Robotics and Computer-Integrated Manufacturing*, vol. 1, no. 1, pp. 65-71, 1984.
51. Siegel, D. M., Garabieta, I., Hollerbach, J. M., "An Integrated Tactile and Thermal Sensor," *Proc. IEEE International Conference on Robotics and Automation*, San Francisco, pp. 1286-1291, 1986.
52. Siegel, D. M., Simmons, L., "A Thermal Based Sensor System," *Proc. SME Sensors '85 Conference*, MS85-995, Detroit, 1985.
53. Siegel, D. M., Narasimhan, S., Kriegman, D. J., Gerpheide, G. E., "The Computational Architecture of the Utah-MIT Hand," *Proc. IEEE International Conference on Robotics and Automation*, pp. 918-924, St. Louis, 1985.
54. Snyder, W. E., St. Clair, J., "Conductive Elastomers as Sensor for Industrial Parts Handling Equipment," *IEEE Transactions on Instrumentation and Measurement*, IM-27, pp. 94-99, 1978.

-
55. Tamai, T., "Electrical Properties of Conductive Elastomer as Electrical Contact Material," *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, CHMT-5, pp. 56-61, 1982.
 56. Ten Grotenhuis, R. L., Moore, T. N., "Development of a Carbon Fibre Based Tactile Sensor," *Proc. SME Sensors '85 Conference*, MS85-994, Detroit, 1985.
 57. Timoshenko, S., Goodier, J. N., "Theory of Elasticity", McGraw-Hill, 1951.
 58. Weast, R. C., Editor, "Handbook of Chemistry and Physics," The Chemical Rubber Company, Cleveland, OH, 1968.
 59. White, R. M., King, A. A., "Tactile Array for Robotics Employing a Rubbery Skin and a Solid-State Optical Sensor," *Proc. IEEE International Conference on Solid-State Sensors and Actuators*, pp. 18-21, Philadelphia, 1985.
 60. Wong, K., Van der Spiegel, J., "A Shielded Piezoresistive Tactile Sensor Array," *Proc. IEEE International Conference on Solid-State Sensors and Actuators*, pp. 26-29, Philadelphia, 1985.